HP Professional

THE MAGAZINE FOR THE BUSINESS & PRACTICE OF HEWLETT-PACKARD COMPUTING

September 1987 ■ VOL. 1, NO. 4

- Database Performance
- Relational Technology
- Las Vegas Hosts INTEREX

DEVELOPMENT

Reusing Code on the 3000

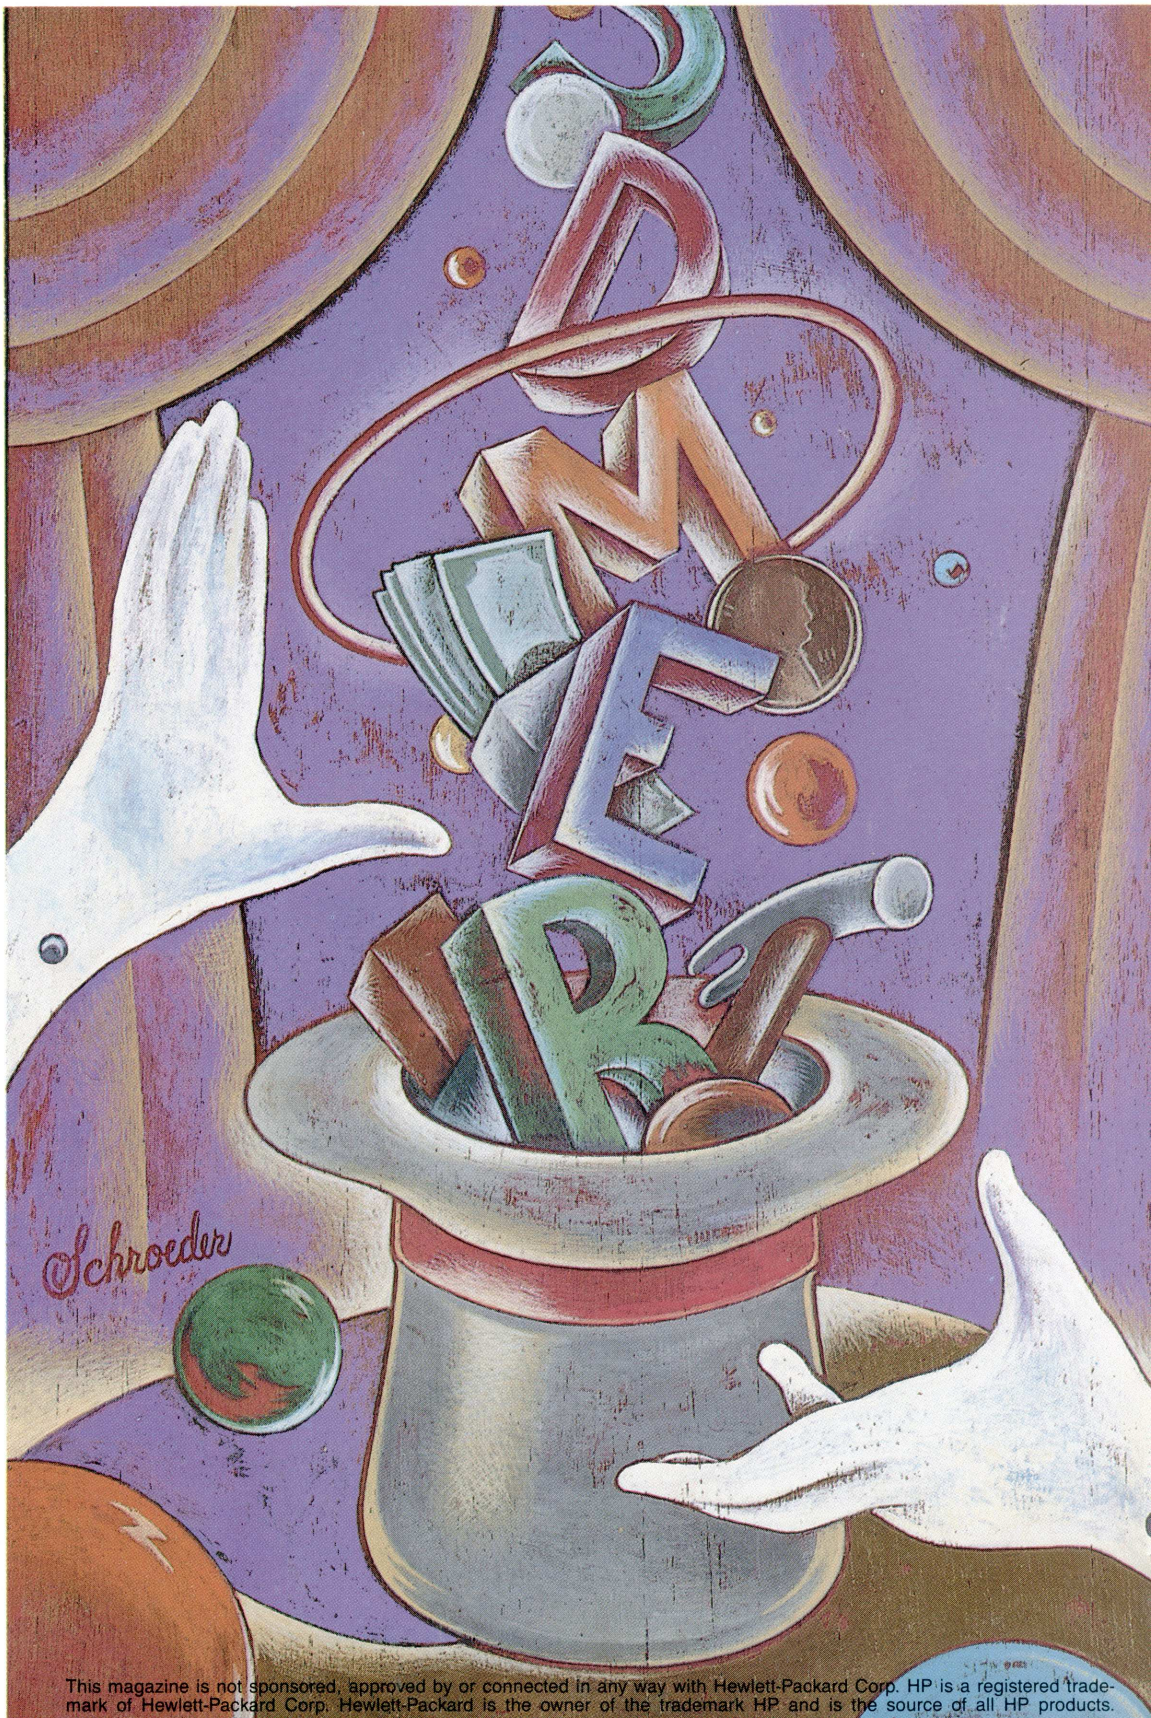

Imagine linking up with a partner that's good looking, intelligent, compatible and affordable to boot...

You imagined it. We created it. The Zentec 8392 Terminal.

When it comes to finding an ideal partner, it's important to do a little comparison shopping—especially if your current partner operates under the name of HP 2392A.

The Zentec 8392 offers everything the 2392A user needs and more. For less. Like a 14-inch diagonal screen. Yet hidden behind its sleek, sophisticated exterior lies a powerful collection of hardware and software specifically designed to outperform other compatibles.

When you work with the Zentec 8392, we promise you a partner who operates in both 80 and 132 columns with a memory better than your own (up to 8 pages). This partner's bright (display available in green or amber), flexible and reliable. Plus, there's a one-year warranty and maintenance contracts available through Dow Jones Service Company.

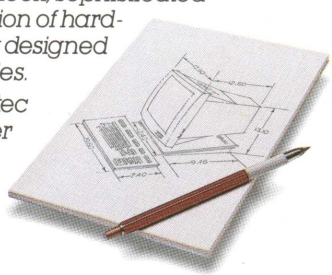

So if you've been playing the silent partner, merely taking what you get because it's there, shop around. Look at Zentec. Once you do, you won't want to look back. For more information, call 408-727-7662 (inside California), 800-332-5631 (outside California), or 011-44-4867-80666 (in the UK). Zentec Corporation, 2400 Walsh Avenue, Santa Clara, CA 95051-1374.

ZENTEC™

THE CUSTOM TERMINALS COMPANY

See us at Interex Booths #505, #507, #604, #606.

ENTER 111 ON READER CARD

Announcing the end of an error.

NSD announces no more hidden production errors. No more wasted effort reviewing every $STDLIST at the end of a run. And no more job listings buried in printouts. Now there's a simple way to save your time and your money.

Our new improved version of the SMP program is here. It's called JobRescue.™ As soon as you try it, you'll understand why. JobRescue runs in the background and continuously examines $STDLISTs. If you choose, it notifies you of errors by displaying a warning on the console. You can then send this warning to other users.

With JobRescue, there is a lot more at your command. For example, now you can keep your $STDLISTs on line for more than 24 hours. Let's say you come in on Tuesday. You can review Monday's reports or hold them for another day. JobRescue lets you hold your $STDLISTS with their flagged errors for up to 127 days!

What's more, you can now specify which $STDLISTs and which errors on those $STDLISTs you want to examine. You can store $STDLISTs without errors in a file or purge them completely. Users can specify that certain $STDLISTs should be printed even if they have no errors.

With all that JobRescue does, it might surprise you how easy it is to use. You simply STREAM a job. In many cases, online commands are only one word long. In fact, because of your familiarity with HP programs, you already know most of the commands. And we guarantee JobRescue for one full year.

To show you just how well JobRescue performs, we're offering a free trial run. The end of an error means the beginning of smooth operations for you.

See you at Interex Booth #629.

ENTER 103 ON READER CARD

NSD

Looking for Trouble
HP 3000 Software

JobRescue is a trademark of NSD, Inc.

Free Trial Run!
800.538.3818, (In CA) 800.538.9058
Canada 800.445.3818, Local 415.573.5923
NSD, Inc. 1670 South Amphlett Blvd.
Suite 103, San Mateo, California 94402

THE DESIGNER PRINTER.

TDP and the LaserJet with PSP / Plus.

If you use TDP, you probably fall into one of two classes: you either have a laser printer, or you want one. If you already have a laser printer, you're probably used to producing high-quality illustrated documents at a whim. If you don't, you're probably working with a daisy-wheel printer regarded by its users as a faithful but aging pet spaniel.

PSP / Plus and the LaserJet can provide you with all the benefits of laser printing at about a quarter the price of the HP2688: full graphics capability, up to 31 fonts in a document, forms, borders, logos — in fact, almost everything that TDP can make a big laser printer do. It works with popular "front end" processors, too, as well as the TDP Mailer facility. In fact, with PSP / Plus, the LaserJet can be used with your own applications and with HP software: HPDraw, DSG, HPEasyChart, HPWord, IDS and IFS all work with the LaserJet, thanks to PSP / Plus.

So if you want to replace the aging daisy-wheel printer, or get to all the power that TDP *really* has, or just want to avoid walking four blocks to the data center's HP2680, try a LaserJet: with PSP / Plus, the LaserJet has joined the HP3000 family.

See us at Interex Booth #105.

299 W. Foothill Blvd. • Suite 230 • Upland, CA 91786
714/985-1581 • 800/858-4507 • Telex 532678 OPT INC

ENTER 113 ON READER CARD

HP2680, HP2688, TDP, HPWord, DSG, HPEasyChart, IDS, IFS, and the LaserJet are products of Hewlett-Packard Company.

OPT

CONTENTS

SEPTEMBER 1987 VOL. 1, NO. 4

28 ***FOCUS:* Issues in Relational Technology**
by Paul Bass
Considerations of transportability, training and performance.

32 ***FOCUS:* The IPC Message-File Feature**
by Gilles Schipper
A TurboIMAGE facility to speed programming and execution.

40 ***FOCUS:* Do Migrating Secondaries Give You Migraines?**
by F. Alfredo Rego
Understanding IMAGE internal techniques.

44 ***FOCUS:* Maximizing Database Performance**
by Mark S. Trasko
Principles of optimized indexing, buffering and data ordering.

50 ***FOCUS:* Something Old, Something New**
by Don Person
Getting more for your database dollar.

54 **A Winning Bet**
HP PRO Staff
A guide to the North American 3000 User Conference and Las Vegas.

58 **Spectrum: Cost vs. Benefits**
HP PRO Interview
An interview with Mike Casteel.

FOCUS

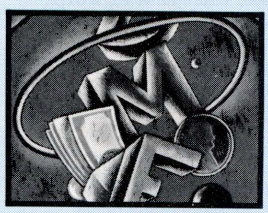

DATABASES/DBMS: Strategies for management of an important corporate asset.

On The Cover:
This month's cover illustration by Michael Schroeder, Reading, PA.

COLUMNS

LAB REPORT: IEEE-488 Controllers For Vectra
by Don Person
A tale of the smart and smarter.....64

DEVELOPMENT: Don't Reinvent The Wheel
by Lisa Burns Hartman
Reusing code on the HP 3000.....74

BUSINESS: Confessions of an HP Horse Trader
by Don Person
Tips on purchasing previously owned, reasonably priced HP computers.....80

OPINION: Computer Fraud Upon Our Children
by Cliff Lazar
LOGO: Is it cheating our children?.....82

DEPARTMENTS

Editorial 10
Letters 14
News & Trends 16
Industry Watch 20
New Products 24
Product Showcase 96
Advertisers Index 98

HP Professional

Publisher: Carl B. Marbach
Editorial Director: R.D. Mallery

Editorial

MANAGING EDITOR Lonni Wright
COPY EDITOR Karen Detwiler
CONTRIBUTING EDITORS Jim Dowling, Lisa Burns Hartman, Don Person
CONTRIBUTORS Paul Bass, Cliff Lazar F. Alfredo Rego, Gilles Schipper, Mark S. Trasko

Design

DESIGN/PRODUCTION DIRECTOR Leslie A. Caruso
DESIGN/PRODUCTION ASSOC. Ruth Ann Leiby
ART/PRODUCTION ASSOC. Timothy M. Kraft
ART ASSISTANT Sue Ann Rainey
PRODUCTION ARTIST Richard Kortz
TYPESETTING/PRODUCTION MaryEllen Springer
DESIGN CONSULTANT Greg Paul, Brady & Paul Communications

Circulation & Administration

VICE PRESIDENT Peg Leiby
CIRCULATION DIRECTOR Mary Wardlaw
CIRCULATION MANAGER Karen Savory
CIRCULATION FULFILLMENT Ruth Henderson, Claire Hollister, Margie Pitrone, Cynthia Reeder, Donna Schmidt
ACCOUNTING Andrea Beneke
COMPUTER SYSTEMS Bonnie Auclair, Kevin Kennelly, Ruth Mermelstein Rory Mulvey

PROFESSIONAL PRESS, INC.

For information on how to contact your sales representative, see page 98. Editorial, advertising sales and executive offices at 921 Bethlehem Pike, Spring House, PA 19477 ■ (215) 542-7008 TWX 910 333 9522 ■ Easylink 62805174 FAX (215) 628-2845

HP PROFESSIONAL is published monthly by Professional Press, Inc., 921 Bethlehem Pike, Spring House PA 19477. Subscriptions are complimentary for qualified U.S. and Canadian sites. Single copy price, including postage, $4. One year subscription rate $30 U.S. and Canada; $60 foreign. All orders must be prepaid. Second class permit applied for and pending at North Wales, PA and additional entry. POSTMASTER: Send all correspondence and address changes to HP PROFESSIONAL, P.O. Box 445, Spring House, PA 19477.

BPA Membership applied for 4/87.

Now you can get more performance with bigger savings in a smaller space.

Hewlett-Packard's new mass storage products deliver more of what you're asking for. More performance. More energy savings. More space savings. And that adds up to lower cost of ownership. In fact, typical payback on our new mass storage products is less than three years.

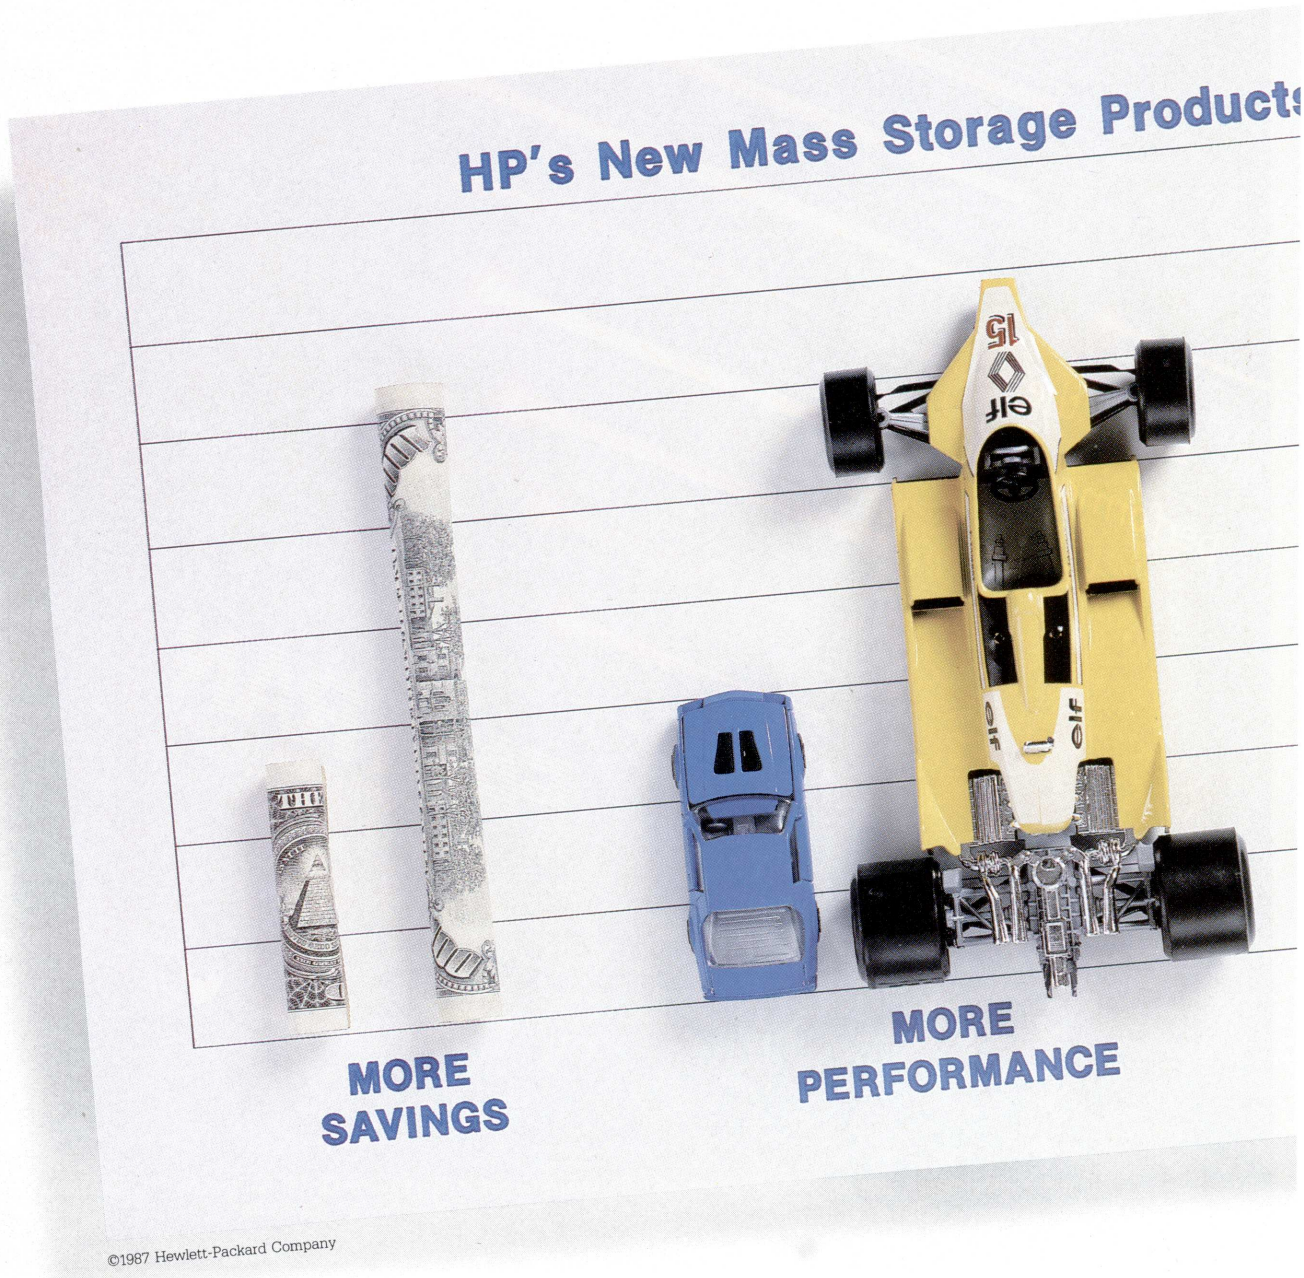

HP's new mass storage units feature a reliable space-saving design that cuts your maintenance costs in half.

Consider our new disc drives, the HP 7937H and 7937XP. When compared to the HP 7933H, the new 7937XP operates more than two and a half times faster while occupying only one-third the floorspace. Plus our new disc drives reduce your monthly computer utility bills by a hefty 80%. And their exceptional reliability cuts your maintenance costs in half. These savings make our new disc drives among the lowest cost per megabyte in the industry.

We give you more in tape drives, too. Our new HP 7980A operates at a fast 125 ips. It rewinds a 2,400-foot reel in 90 seconds flat. And its reliable design cuts your monthly maintenance costs by more than 50%. The HP 7980A features a sleek new casing that mounts horizontally into a 19-inch wide rack. So there's enough room left to add another drive on the bottom and double your capacity. Its new autoload feature makes changing tapes faster and easier than ever to save your operator valuable time.

Now with HP's trade-in programs you can get even more for your money. Call your local HP sales office or dial 1-800-752-0900, ext. 906A and ask for the HP More Pack.

we never stop asking "What if..."

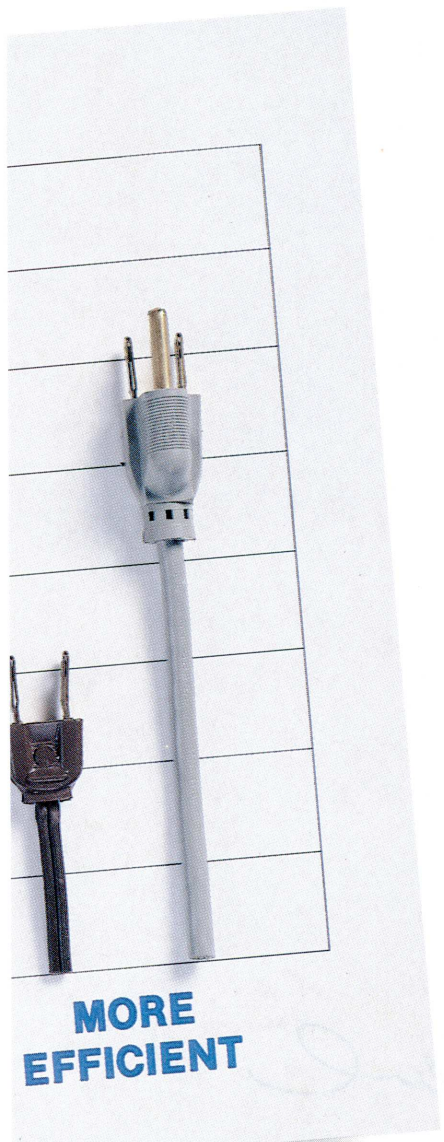

HEWLETT PACKARD

ENTER 200 ON READER CARD

It's Time To Grow

MPE/XL, the Spectrum operating system, has to begin to distance itself from MPE because it runs on the older 16-bit HP 3000 computers. Old ideas die hard, but now lessons must be learned as we move towards the 1990s.

An example of old thinking is the oft-mentioned type-ahead, or lack thereof. MPE doesn't have a type-ahead feature; MPE/XL doesn't either. The current thinking at HP is that only word processing programs need this and they have addressed word processing with attached PCs and workstations. This shows a distinct lack of vision in operating system design.

Type-ahead buffering can be used by many other applications besides word processing. Specifically, it can increase an operator's efficiency, for any type of input, by not having the operator wait for the computer. Just watch any experienced computerist and he almost always will be ahead of the computer.

Type-ahead is only one stumbling block. File structures and organization, database architecture, networking constructs and command languages are a few others that need to be addressed. The Spectrum computer line will address more memory, read from more/bigger disks, communicate with other computers and serve more users than the older 3000 series. Spectrums will be called on to do more, sometimes doubling as a technical computer as well as a commercial one. These needs and power will require a more robust operating system than is currently available.

We can expect the Spectrum line to span a range of computing that will be impressive — from desktop to corporate mainframe. The operating system needs to be simple enough for a single user workstation and powerful enough to serve thousands of users in a corporate environment. Like the hardware, the operating system needs to be stretchable.

While compatibility is an important issue, we should remember that for the Spectrum and older 3000s it should be UPWARD only. Things that work on a 16-bit 3000 should work on the newer Spectrum, but there *must* be things on the Spectrum that won't work on the traditional computers.

Developing new constructs, tools and functionality must be a high priority at HP. What will set this computer system apart from the competition will be the software.

Fortunately, the greater address space and architecture of the Spectrum will allow the operating system engineers wide latitudes in new functionality. Importantly, *they can't think of them all.* They need your help. That is what a users group is for. It's time for us to make clear to HP what WE want, not what they want to give us. If you WANT type-ahead, demand it. HP will listen.

While sometimes late and often too slow, HP has responded to the needs of its loyal customer base by fulfilling the promise of Spectrum. Spectrum needs not only its new hardware architecture for future success, but also an operating system worthy of this fine architecture.

IMACS *DataExpress* SERIES

The Ultimate Downloading Machine

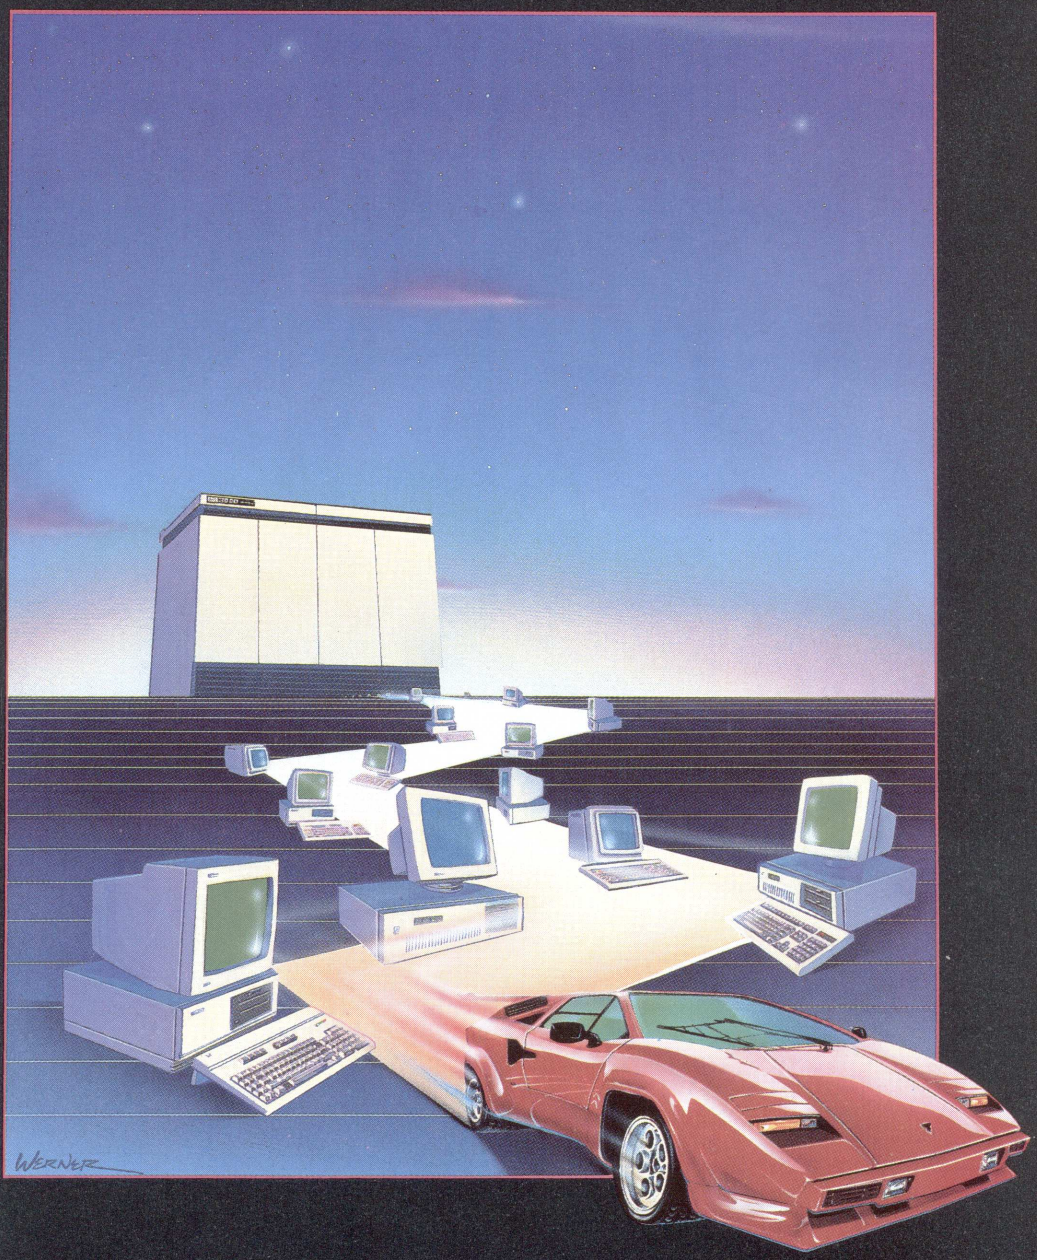

IMACS DataExpress integrates the PC Workstation with HP-3000 applications

If you are thinking download, study the options, then choose power, performance and classic quality

Data extraction and download have never been easier
Call IMACS and ask about the DataExpress Series

See us at Interex Booth #327.

IMACS®

IMACS SYSTEMS CORPORATION
4676 Admiralty Way, Suite 516
Marina Del Rey, CA 90292
(213) 823-4577

ENTER 219 ON READER CARD

United States	United Kingdom	Netherlands	Canada
(800) 445-9519	(01) 541-0242	(010) 139310	(403) 259-2478

IF YOU HAVE MORE THAN ONE

BACK

BACKCHAT GIVES YOU A DISTRIBUTED DATA BASE

BACKCHAT logically links multiple IMAGE data bases in a multi-processor HP3000 network. It simplifies the design of distributed processing systems and ensures that failures in the network have minimal impact on your operations. Data can easily be shared between computers and applications. BACKCHAT is the only way to spread an IMAGE data base over more than one HP3000 without ad-hoc processing.

BACKCHAT FOR A FAULT TOLERANT HP3000
BACKCHAT can mirror one or more data bases onto a second HP3000 in real time,

IP3000 YOU NEED BACKCHAT

CKCHAT™

giving you a faulty tolerant configuration. So if your primary system fails you can switch over to a back-up system in seconds. BACKCHAT offers more than SILHOUETTE but costs less.

BACKCHAT FOR IMPROVED PERFORMANCE

Because you can off-load processing to a second machine you can improve application system response with BACKCHAT.

See us at Interex Booths #128, #130.

FOR MORE INFORMATION CONTACT:

PROACTIVE SYSTEMS

5053 LA MART DRIVE, SUITE 101, RIVERSIDE, CALIFORNIA 92507 TEL: 714-683-7921

PO BOX 7102, BLOOMFIELD HILLS, MICHIGAN 48302 TEL: 313-333- 7200

CENTRAL COURT, KNOLL RISE, ORPINGTON, KENT BR6 0JA, ENGLAND TEL: (0)689-77933 TELEX: 9413362

PROACTIVE SYSTEMS

INFORMATION SYSTEM CONSULTANTS AND SOFTWARE SUPPLIERS

ENTER 240 ON READER CARD

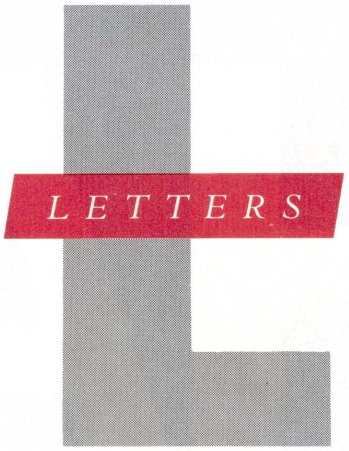

IMAGE Performance Myths

I found the article by David Merit, "IMAGE Performance Myths," quite interesting. Among all the topics that came to my mind, I would like to share a few with your readers:

David's mission is to "squelch the 'Half-Truths' About HP's Efficient DBMS." This is certainly a worthy cause and he has my support (and the support of thousands of other serious students of IMAGE).

David says "[in a detail dataset] IMAGE reads only to the high-water mark (which reflects the highest active entry)." The first half of the statement is true, but the second half (in parenthesis) is, to paraphrase David, "the downfall of the statement." These are the facts: The high-water mark may or may not reflect the highest active entry.

Try this experiment: Take a DETAIL dataset whose high-water mark is at 0 (for instance, use a database that you have just created with DBUTIL). Add some entries with QUERY to this detail (call them entries 1, 2, 3, and 4). So far, so good: the high-water mark is, indeed, at the highest active detail entry (number 4). Now, delete entries 1 and 4; do a QUERY "find all" command, followed by a "numbers" command, to see the entry numbers of the active entries (which are entries 2 and 3). Exit QUERY and run DBCHECK (or TBCHECK) on the detail in question to get the report on the high-water mark. You will see that the high-water mark (at entry 4, as reported by DBCHECK) refers to the highest-ever active entry, not to the highest active entry (which is entry 3, as reported by QUERY's "numbers" command).

The difference between "highest-ever" and "highest" is the difference between life and death when you deal with the privileged internals of IMAGE (or TurboIMAGE) structures.

Regarding primary paths, David says, "It's also possible that a primary path will move. Beware of the technique in some database utilities of deleting a path and adding it back in to correct a broken chain by rebuilding them all."

This is an obvious reference to the good-old Adager method which I introduced back in 1978 as a pioneering technology and has been used successfully by thousands of Adager users all over the world for almost a decade. I must stress the fact that the old Adager allowed you to redefine the primary path, if you so desired. Therefore, there is nothing to "beware" except for the inconvenience of having to perform two Adager functions. To simplify the user's task, in 1986 I introduced yet another pioneering technology which made obsolete the Adager technique mentioned by David: Adager's PathFix function rebuilds all the chains in a path without having to delete or add the path.

Let's get back to primary paths. David says, "The assignment of the primary path is significant, because the detail entries are unloaded in the order of the chains on the primary path. Since the entries are reloaded in that order, the primary path dictates in what sequence the entries physically will reside following the reorganization."

Fortunately, Adager made database reloads obsolete back in 1978. (Are you up to date?) To appreciate Adager's approach, try this: Use Adager's DetPack function to repack a detail dataset along any path of your choice, whether primary or not. As a standard (and exclusive) Adager bonus, you will preserve the current chronology of all the chains in all paths.

F. Alfredo Rego
Adager
Antigua, Guatemala

Address letters to the editor to the *HP PROFESSIONAL* magazine, P.O. Box 445, Spring House, PA 19477-0445. Letters should include the writer's full name, address and daytime telephone number. Letters may be edited for purposes of clarity or space.

A disk drive that three can enjoy.

Hewlett-Packard users, drink up.

If the high cost of storage and networking has left you stranded, you're in for a treat.

Introducing MultiPort. The HP-compatible disk drive that can be shared by either two or three computer users.

MultiPort is one complete, cost-effective system for both storage and backup. That's because each user is allotted one *equal* amount of storage on the fixed disk. In addition, they also have access to a 20MB removable Winchester cartridge, and a Shared Information Space (SIS) that serves as a networking alternative.

So it's like having three (or two) separate hard disks and a backup unit on one space-saving system.

Share and share alike.

If you want to share files from one user to another, you just download your data onto the removable cartridge (just like you would a floppy), remove the cartridge, pop it into your drive and away you go.

There's no problem of accidental data address. And because the 20MB cartridges are removable, you have unlimited storage capacity. You have unlimited security as well, since you can lock the cartridges up.

SIS–The refreshing difference.

MultiPort's unique Shared Information Space is also available for low-cost networking. When the SIS is configured, it acts like an extra, "separate disk" that everyone can share.

Thus, everyone can access it and work off the same file, but only one at a time. Security and traffic features prevent simultaneous access of the SIS "disk."

Storage choices that satisfy.

MultiPort is available in 20, 40, 50 or 70MB models. Software is included. The system is compatible with HP9000 Series 200/300/500 computers and fits HP's ITF footprint specifications for easy integration into your HP environment. Choices include:

8220: 20MB fixed & 20MB removable
8240: 40MB fixed & 20MB removable
8250: 50MB fixed & 20MB removable
8270: 70MB fixed & 20MB removable

For more information on all our MultiPort systems, call or write today. Bering Industries, Inc., 280 Technology Circle, Scotts Valley, CA 95066, 408-438-8779. Inside California, call **800 533-DISK.** Outside California, call **800 BERING 1.**

BERING™
A MOUNTAIN COMPUTER, INC. COMPANY

Innovative Storage for Hewlett-Packard.

See us at Interex Booth #622.

ENTER 196 ON READER CARD

HP Addresses New Markets With DEC, IBM, ASCII Terminals

Also Adds New Terminals For HP Systems

Expanding its commercial-workstation family, Hewlett-Packard recently introduced new low-priced, full-featured terminals for IBM, DEC and general-purpose ASCII environments. The company also replaced two existing terminals for its own computer systems with new models that offer more features at a lower price.

This is the first time that HP has introduced terminals specifically designed to run on other vendors' systems. Called the HP 700 display-terminal family, the new alphanumeric terminals are designed for business, technical and manufacturing applications.

List prices on the HP terminals for ASCII, DEC and IBM environments are five, 28 and 45 percent lower, respectively, than those offered by current leaders in each market.

According to forecasts by Dataquest Inc., a market-research organization, annual unit shipments of terminals are expected to grow from 2.6 million in 1987 to nearly 3.6 million in 1991 — 38 percent during the five-year period — as customers replace older equipment.

All terminals have 14-inch diagonal screens and are engineered to produce crisp, clear characters. For example, the use of a "half-dot" shift creates smoother curves in characters. Each terminal also includes standard ergonomic features such as tilt-and-swivel monitors and adjustable keyboards.

Terminal screens are available in green and amber. Additionally, the new HP and DEC-compatible terminals are available in soft white, in which black letters are displayed on a white screen. All controls — such as power, brightness and contrast — are located on the front of the terminal for easy access.

All terminals include HP's one-year warranty.

The HP 700/41 display terminal ($375) is the lowest-priced member of HP's new terminal family. This entry-level, general purpose ASCII terminal works with a variety of computer systems and software. Users may select from an on-screen menu for compatibility with terminals by Wyse, TeleVideo, ADDS, Lear Siegler, Qume and Hazeltine.

The full-featured keyboard includes 58 programmable keys, with 16 function keys that each serve a second function by pressing the shift key.

The HP 700/22 display terminal ($575) works with programs designed for the DEC VT220, VT100 and VT52 terminals and other applications that follow the ANSI 3.64 protocol. HP's terminal for DEC applications has four pages of memory for the display — four times that of DEC's VT220 terminal — which allows more information to be displayed by scrolling up and down.

The keyboard has 106 keys in the same layout as DEC's keyboard, but includes twice the number of programmable keys.

The HP 700/71 display terminal ($695) can run applications requiring an IBM 3191 Model A or B terminal in a 3270 Information Display Systems environment.

In addition to the terminal controls for power and brightness, a security lock is located in the front.

HP offers 122-key and 102-key keyboards in the same arrangement as those from IBM.

The HP 700/92 ($895) and HP 700/94 ($1,095) terminals for HP computing environments offer more features and a lower price than the HP 2392A and HP 2394A terminals that they replace.

Added features for the HP 700/92 terminal include 80- and 132-column viewing modes, storage of up to eight pages of display memory, a second printer port and eight function keys that can be doubled by pressing the shift key.

In addition to the features of the HP 700/92, the HP 700/94 includes forms cache, which allows an average of 25 forms to be stored in the terminal's memory. In addition, 11 data-edit checks, which are used to reduce errors, are also done at the terminal rather than on the host. Both features help streamline the data-entry process by reducing the need for interaction with the host computer.

The HP 700/92 and /94, which are used with HP's commercial and technical systems, take advantage of block-mode communications in HP 3000 software packages.

Portable PC Family Expanded

New Vectra Designed For Office, Sales Pros

The first two members of a new family of industry-standard portable personal computers were introduced recently by HP.

The HP Portable Vectra CS and the HP Portable Vectra CS Model 20 computers are for office and sales professionals who need the functionality of a desktop computer and the convenience of a battery-powered portable. The portables can be used on the road, as well as be converted into desktop PCs, which are adaptable to external monitors.

With this announcement, HP now offers two types of portable computers: The HP Portable Plus, available since 1985, will continue to be offered for sales and service professionals who need a lightweight portable with built-in ROM (read-only memory) software. The new portable PCs, based on an Intel 8086-compatible CMOS microprocessor, address business-professional needs for a portable computer with industry compatibility, increased storage and greater functionality.

The HP Portable Vectra CS personal computer is the first battery-powered portable PC to offer four internal I/O expansion slots, making it adaptable to different customer needs and future technology. Users may add adapters for increased memory, communication devices, or peripherals such as printers and plotters. These adapters can be added without increasing computer size.

This new PC provides two 1.44-MB flexible-disk drives, while the HP Portable Vectra CS Model 20 PC has a 20-MB hard disk and one 1.44-MB flexible-disk drive. The high-density, 3½-inch disk drives offer twice the average storage of other portables and are compatible with IBM Personal System/2 disks as well as industry-standard 720-KB disks.

Both systems come standard with 640 KB of user memory. Up to 6 MB of EMS RAM can be added to the Portable CS, and up to 4 MB to the hard-disk model.

The 12-inch diagonal liquid-crystal display (LCD) is removable, so that the system can be connected to an external monitor for use as a small-footprint desktop PC. A standard display/printer adapter can drive either the LCD or an external color or monochrome monitor. The LCD, based on "supertwist" technology for high contrast, offers 640 by 400 resolution and can be adjusted to any angle.

A full-size keyboard, compatible with the IBM PC/AT and PS/2 keyboards, features 92 keys, including 12 function keys and a separate numeric keypad.

A single battery charge runs the HP Portable Vectra CS for up to 10 hours and the hard-disk model for up to four hours. A battery "fuel" gauge indicates the remaining charge and gives a low-battery warning, allowing ample time to back up files.

An HP Portable Vectra CS with two flexible-disk drives is $2,495 and is available now. The HP Portable Vectra CS Model 20 PC with one 20-MB hard disk drive and one 1.44-MB flexible-disk drive is $3,595 and is scheduled to be available in the first quarter of 1988.

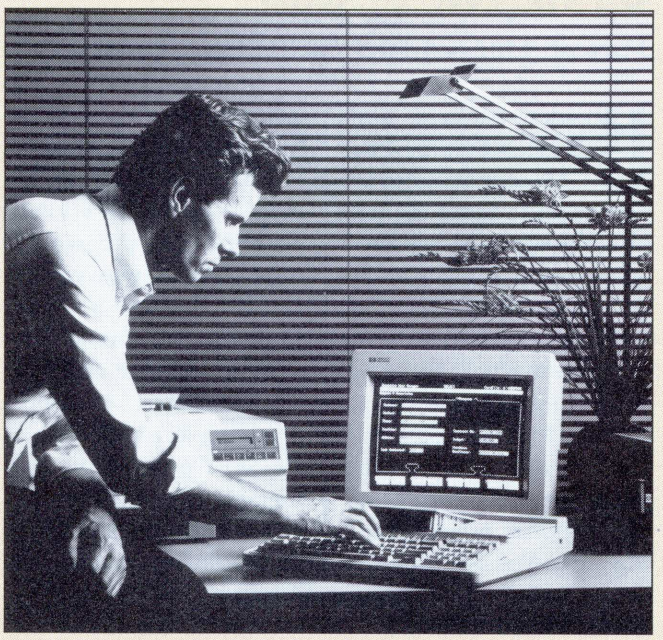

The portable Vectra CS's LCD can be replaced with a full-size color monitor for use as a desktop PC.

Bentley College Purchases 2,200 Portable Vectra CS Computers

Freshmen, Sophomores Issued Portables

Bentley College, the largest undergraduate college of business in New England, has purchased 2,200 HP Portable Vectra CS personal computers for its students and faculty.

Incoming freshmen and sophomores at the Waltham, MA, institution will use the portable PCs for such projects as creating spreadsheets for finance courses, solving problems in accounting and writing class papers. Students can use the portable computers in their dormitory rooms, at study sessions or away from campus on weekends and vacations.

Each member of Bentley's full-time faculty is issued either a desktop or a portable HP Vectra PC to develop and distribute class assignments, prepare lecture notes and perform research. Because the new portables

are compatible with industry standards, professors can make assignments using

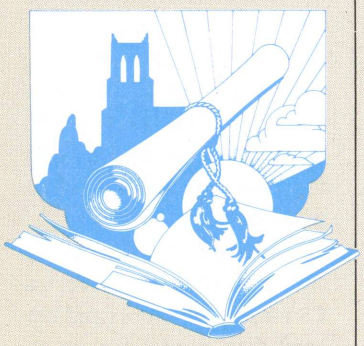

textbook software without modifying the programs.

The college, a national leader in integrating computers into the curriculum, began a pilot program three years ago that equipped selected students with the HP Portable Plus laptop computer.

The success of that program prompted the college to issue portable computers to all freshmen and sophomores. This practice will continue with the more powerful HP Portable Vectra CS, which supports standard educational software.

Communication ports in dormitories give students better access to on-campus minicomputers and HP Vectra PCs. A word processing package, spreadsheet program and BASIC programming language will be used with the HP Portable Vectra CS computers.

HP To Host VAB Conference

Focus On Technical Systems

Hewlett-Packard will offer its third annual conference for Value Added Businesses September 23-25 in Steamboat Springs, CO. The conference is designed for all HP technical third-party software suppliers, resellers, system integrators and OEMs. Value Added Businesses that are interested in working with HP's technical computers — the HP 1000 and 9000 manufacturing, engineering and general purpose personal to superminicomputers — will attend.

The Value Added Business (VAB) Conference will feature business workshops, presentations by industry business consultants, overviews of HP corporate strategy, program updates and HP product reviews. The list of invitees to the Conference will come from current participants in the HP Plus Program (HP's Value Added Software Supplier Program) and prospective third-party vendors as identified by the HP Field Sales Force.

Hosted by HP's Technical Systems Sector, the conference will be built around the theme of "Sharing Strategies for Success." Participants will be encouraged to share their experiences and marketing plans while selected members of HP's management team will be present to discuss their business plans and product strategies.

The conference also will supply workshops that will help the VABs be successful in general as well as how to work with HP. Topics include: Improving Software Productivity, Market Segmentation, Financial Management for Growth, Cross Licensing and Working With the HP End-User Sales Force. Workshops will be led by HP personnel and industry consulatants and will be designed to encourage interaction and dialog.

Noted industry consultants, Mack Hanan and David Flegal, will present special workshops at the conference. Hanan, the inventor of "Consultative Selling," will present a half-day version of his workshop on how to establish the "consultative" sales relationship with customers. Flegal's presentation will focus on strategic planning and how it can be applied in the software business.

Agreement Puts Smalltalk-80 On HP 9000 Systems

Part of Value-Added Business Program

Hewlett-Packard and ParcPlace Systems (Palo Alto, CA) recently signed an agreement to make the full line of Smalltalk-80 software-development systems available on the HP 9000 series of technical workstations.

The Smalltalk system, a product of the research at Xerox's Palo Alto Research Center (PARC), is known for its abilities to manage complex software projects, rapidly prototype new products and integrate the accessibility of other applications.

The Smalltalk agreement is part of HP's value-added business program for third-party software vendors in which HP provides marketing and porting assistance to software suppliers.

ParcPlace, currently operating as a subsidiary of Xerox Corporation, was founded to market, support and continue development of the Smalltalk-80 environment.

New IEEE Standards Group To Study 10-Mbps StarLAN Standard

HP Makes Formal Request

A new 10-Mbps StarLAN local area network standard that will extend the capabilities of telephone wiring to meet future office networking requirements will be explored by a new IEEE standards study group.

This was the outcome of a formal request made by Hewlett-Packard to the IEEE 802.3 (CSMA/CD) Standards Committee.

At a July IEEE committee meeting in Vancouver, British Columbia, HP, with

the support of a dozen other companies, proposed that a study group look into a 10-Mbps version of the current 1-Mbps StarLAN standard that relies on twisted-pair telephone wiring to send information among multivendor PCs and departmental office minicomputers.

HP said the proposed high-speed StarLAN specification will permit even greater use of telephone wiring for computer communications, supporting many more users who will be able to transfer larger amounts of information more efficiently. If accepted as a networking industry standard, the 10-Mbps StarLAN technology will allow companies to expand their current StarLAN multivendor networks, and mix and match the 1-Mbps and 10-Mbps networks to meet varying office communications requirements.

"In the next few years, as higher-performance (386-based) PCs become accepted and widely used, PC owners will be able to take full advantage of the transmission speed of today's StarLAN-standard networks — and some will require the additional bandwidth of tomorrow's 10-Mbps StarLAN-standard networks," said Willem P. Roelandts, General Manager of HP's Information Networks Group.

After taking a close look at these market needs and the lengthy standards process, HP encouraged the IEEE 802.3 committee to study a standard that capitalizes on the evolving technology.

The new IEEE study group, composed of interested companies, began investigating the 10-Mbps specification at an August meeting. The group explored the scope, objectives, applicable technologies and feasibility of a 10-Mbps StarLAN standard effort. It's anticipated that a formal standard-development effort will be initiated within the coming months, HP said.

Template Graphics Signs Agreements With HP

To Develop Custom FIGARO

Template Graphics Software (TGS) and Hewlett Packard (Ft. Collins) jointly announced a formal engineering and marketing agreement for custom versions of TGS's FIGARO software for the high-end HP 320/SRX, 350/SRX and 825/SRX graphics workstations.

FIGARO, TGS's implementation of the Programmers Hierarchical Interactive Graphics System (PHIGS), is a high performance device and computer-independent standard designed for 2D and 3D graphics applications requiring hierarchical data structures, geometric modeling, rapid display modification and interactive input.

TGS is working with HP to develop a tightly integrated implementation of FIGARO. This implementation will utilize the high performance HPSRX graphics accelerator for transforming, clipping and rendering of geometric primitives and hidden surface removal. Using this implementation, FIGARO application programmers will be able to easily develop applications that will perform many times faster than implementations not designed to use the graphics accelerator.

HP joins SUN, Silicon Graphics and Stellar Computer by jointly developing, marketing and supporting high performance implementations of the PHIGS Standard via FIGARO.

FIGARO's internal design defines a high-level Graphics Engine Interface (GEI) which allows it to utilize the high-performance graphics accelerators that support 4x4 floating point matrices and can transform and clip 2D and 3D lines, polygons and text specified in full floating point coordinates. The Graphics Engine Interface (GEI) allows additional workstation manufacturers and system integrators to easily support FIGARO on other platforms.

New European Software Group Formed

Orbit To Market BACKUP/3000

Joerg Groessler GmbH of West Berlin and its marketing arm, Productivity Tools Ltd. of London, have merged all business activities under a new name and corporate structure, the ORBIT group of companies (see chart).

The ORBIT group is marketing its BACKUP/3000 software, a high performance STORE and RESTORE utility, directly and through selected agents. ORBIT will be introducing ONLINE-BACKUP for all series of the HP 3000 later this year and has plans to open a support company in the U.S. at that time.

ORBIT has support facilities in:	
United Kingdom:	ORBIT Software (UK) Ltd., phone (44) 1-948-5166
West Germany:	ORBIT Software GmbH, phone (49) 30-852-7097
Holland, Belgium:	ORBIT Utilities (Benelux) B.V., phone (31) 4744-2214
France:	ORBIT Logiciels Sarl, phone (33) 40-947-405
Spain:	ORBIT Software (Espana) S.A., phone (34) 3-238-3055
Scandinavia:	ORBIT Software (Scandinavia) A.B., phone (46) 303-850-50

Connecting PCs/Workstations

INDUSTRY WATCH

Dave Mallery and Carl Marbach

Every day, it's becoming easier to connect PCs or workstations to larger computers. It's time to ask, "What do I do once I have the connection?"

Terminal emulation allows PCs to act as terminals attached to our main computers. File transfer routines that usually accompany terminal emulation packages allow us to transfer files between the mini and the micro. A problem with some of the software is that it requires too much intervention by the often naive PC user. Anything we can do to insulate the PC functions from the ones that need to be performed on the mini is a real plus.

In a two-pronged assault on the market, WRQ is bringing new products and programs to the PC connectivity market.

As an example of how it can be either simple or complicated is demonstrated by just trying to print a file that exists on a PC to a printer on the mini. First use the file transfer protocols to move the file from the micro to the mini. Then, using the constructs on the mini, print the file. Alternatively, we could have a team of programs, one on the PC and one on the mini, that would print files from the PC to a printer on the mini *without* our intervention.

Another example is interaction between HPDESK, your PC *and* other PC users. Complicated to do by hand, it's simplified with code that exists on both the PC and the 3000.

IN A TWO-PRONGED ASSAULT on the market, WRQ is bringing new products and programs to the PC connectivity market. One is a family of products that "complement" the Reflection Series of terminal emulators. The other is a program for VARs that should create far deeper market penetration.

The Complements family adds functionality for the end user without leaving the familiar Reflection environment. The first member of the family, RSVP, allows the PC user to spool PC printing to any printer attached to the host HP 3000.

The second Complement package assists HPDESK users both in the unattended up and down loading of HPDESK messages and also in the transmission or receipt of files between the Reflection workstation and other HPDESK users.

The VAR program, named ROI (Reflection Office Integration!), is a marketing program designed to assist third-party vertical system houses to propogate Reflection software into new areas. ROI will help software houses add confidence in their new vertical products while avoiding wheel reinvention. Benefits of ROI will include joint advertising, use of the Reflection Logo and marketing assistance.

RSVP (Reflection Spooled Virtual

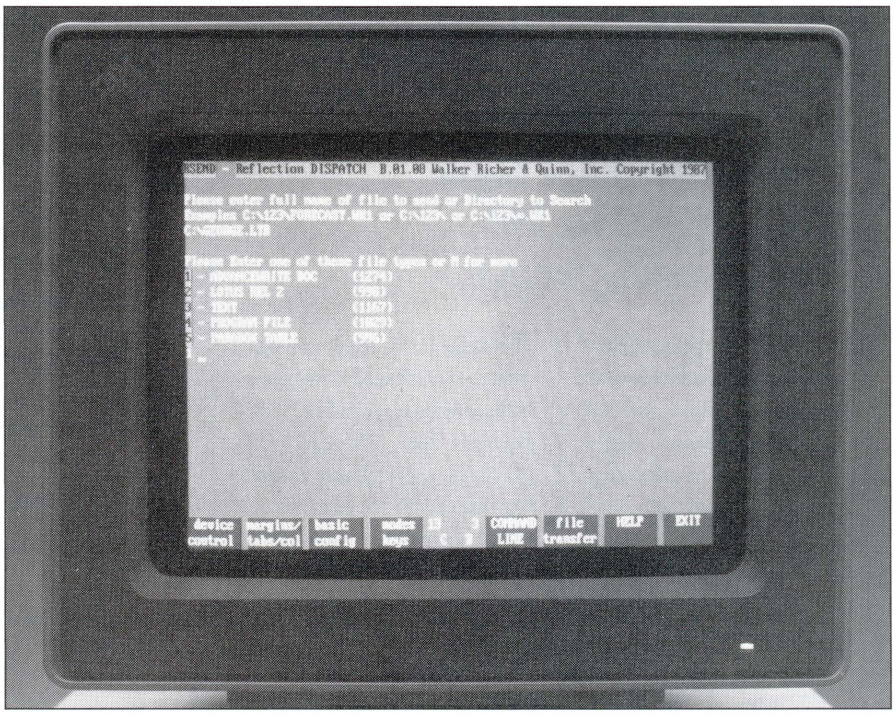

Dispatch lets a PC user connected to host with Reflection send many kinds of PC files through HPDesk Mail to other HPDesk users.

Introducing, the 4GL Environment

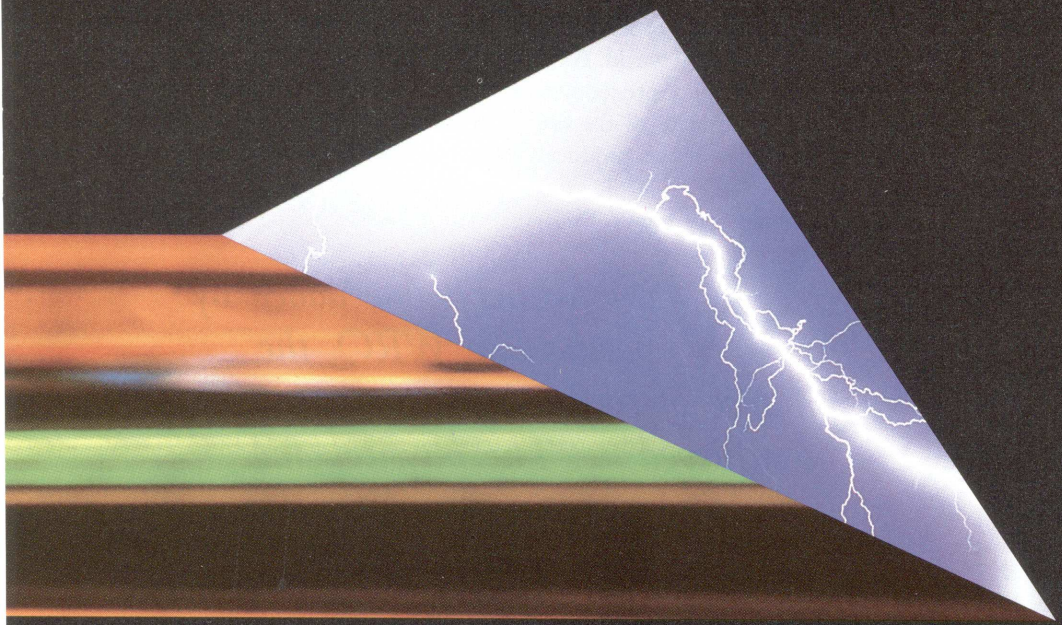

The fourth generation language (4GL) was a revolution. It allowed programmers to code 10-20 times faster than COBOL! That's a benefit that is difficult to live without. However, when developing and running application systems, there is more to life than coding. Much more.

Enter the 4GL Environment. The next revolution.

If you also want to improve the speed and quality of systems analysis and design, produce user manuals automatically, integrate PC's into your applications and still squeeze the most performance out of your hardware.

Then you need to find out about the SPEEDWARE ENVIRONMENT.

Call us . . . . . Right now . . . . . Or else!

Marketing Dept.
Telephone (416) 678-1841

Infocentre Corporation

infocentre

For more information contact:

Infocentre Corporation
7420 Airport Road
Suite 210
Mississauga, Ontario
L4T 4E5

In Europe call:
London (01) 828-1897

See us at Interex Booths #706 - #712.

Only one product does it all,

SPEEDWARE™
ENVIRONMENT

Offices throughout the world including:
New York, Los Angeles, Montreal, Toronto, London, Paris, Sydney.

ENTER 198 ON READER CARD

Printer) has three main components. The first is a filter program resident on the host 3000 that filters non-ASCII characters from the incoming files to be printed. Any type of PC file may be transferred to an HP printer that recognizes ASCII carriage control. The typical use of RSVP is to transfer word processing documents, spreadsheets or PC graphic output to user-selected printers on the HP 3000.

The second piece of RSVP is a memory-resident spooler for the PC. This code traps the printer-directed I/O and redirects it to disk. The spooler retains all the formatting and display enhancement codes supplied by the PC. When the file is subsequently transferred to the host printer, it is identical in appearance to what would have appeared on the local PC printer. The spooler is configured and controlled through a pop-up screen that is always available to the user.

The last component is a Reflection command language program that executes a standard file transfer to the HP 3000, invokes the filter program and finally sends the file to the desired printer.

HPDESK is a standard tool for office automation in the 3000 world. Poor performance and complex procedures hinder the otherwise state-of-the-art features and flexibility of this product. HPDESK users who spend most of their time "off" the 3000 and "on" their PCs can use DISPATCH as a convenient way to compose messages locally and to periodically query their in-tray back on the HP 3000.

Another major challenge for HPDESK users is file transmission to other HPDESK users or to printers on the HP 3000. Dispatch simplifies the procedure by providing a question and answer format that requests source and destination. Word processing documents, database and graphic files are easily transmitted through HPDESK with minimal user involvement.

With further value added by parties like WRQ, the promise of connectivity can be achieved . . . The physical and logical links are in place, it's up to us to use them properly.

Hardware for either product is an IBM PC, HP Vectra, Compaq or other PC compatible with serial or LAN communications, 512K of memory and a hard disk. Software requirements are PC-DOS or MS-DOS 2.0 or higher and Reflection 1, 1+, 3, 3+, 7+ or higher.

CONNECTIVITY IS UPON US. The vendors (PCs/HP) have given us the computers and the tools to physically connect them together. With further value added by parties like WRQ, the promise of connectivity can be achieved. Like most tools, however, it will be the end users who dream up the real applications that make them worthwhile.

In the end, no one remembers the tool that makes it all possible, just the application that does the job and solves the problem. If you're not thinking about how these tools solve your problems, and what the new functionality can do for you, now is the time to start.

The physical and logical links are in place, it's up to us to use them properly.

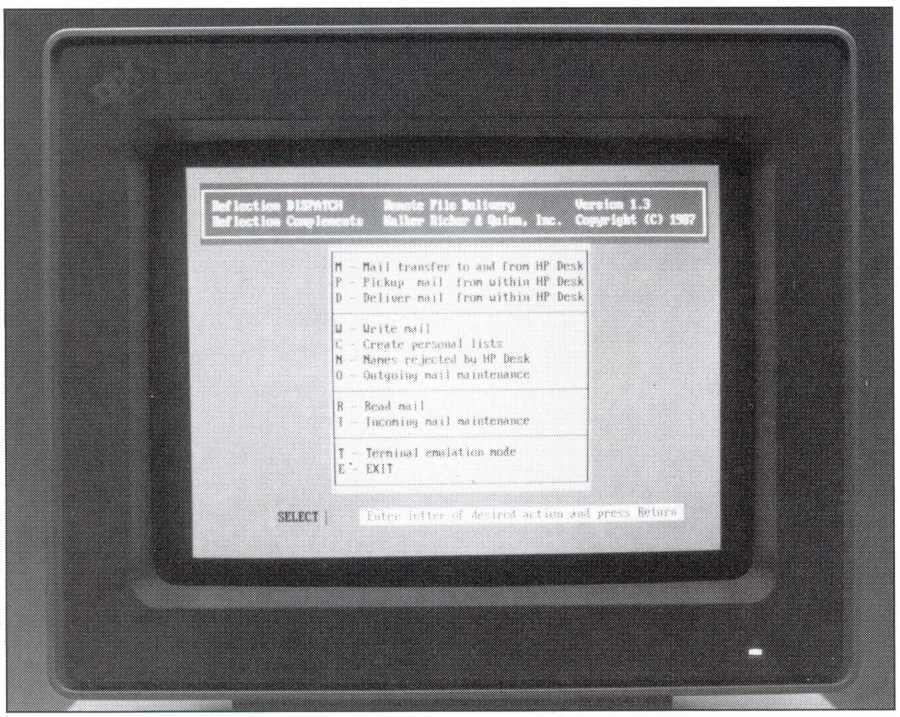

Using Dispatch, PC users connected with Reflection can read and compose HPDesk messages while offline.


See us at Interex Booths #224, #226.
DB GENERAL
Protecting your database
BRADMARK
(713) 621-2808
10
J
Q
K
A
You've got the winning hand!
ENTER 234 ON READER CARD
BRADMARK

NEW PRODUCTS

RDI-3000 Available For 3000 Users

Purvis Systems Incorporated has released its relational database interface, RDI-3000 for HP 3000 users. This software interface provides the handshake between the HP 3000 series hosts and the Intelligent Database Machine (IDM) built by Britton Lee Incorporated, Los Gatos, CA. Now HP 3000 users can extend the capacity of the host by regaining CPU resources previously lost to database and communication software overhead.

The RDI-3000 is a system of computer programs designed to provide a transparent interface for the user. The KEY package, the basic purchase unit, comprises a serial port handler, multiple data stream manager, time run library, interactive parser and database administrator utilities. This interface provides access to the IDM directly from terminal queries or applications software.

The RDI-3000's Run Time Library allows application programs to access the database directly and the Interactive Parser unlocks the relational database power by permitting end users to build hypothetical relations and make ad hoc queries.

Contact Purvis Systems Incorporated, 1272 West Main Road, Middletown, RI 02840; (401) 849-9242.

Enter 900 on reader card

I/O Mate Improves Program Throughput

Program throughput can be improved for programs that read IMAGE databases, MPE flat files, and KSAM files with Running Mate's I/O Mate. I/O Mate requires no source code changes and can be installed into any program within seconds. Chained reads of IMAGE/3000 and TurboIMAGE/3000 datasets can be improved by a factor of two. Serial reads of datasets and also some MPE and KSAM files can show improvements. I/O Mate is compatible with both in-house developed and purchased software including: Query, Quiz, Inform, ASK, Business Report Writer, Collier-Jackson and MCBA.

I/O Mate maintains performance by using no-buff, multirecord I/O and reading into main memory large data chunks from disk. It uses extra data segments to reduce disk I/O and thus decrease CPU resource use and elapsed time of I/O intensive programs.

Contact RunningMate, Division of Pacific Coast Building Products, 3001 I Street, P.O. Box 160488, Sacramento, CA 95816; (916) 325-3663 or (800) 824-9046; in CA, (800) 628-8030.

Enter 901 on reader card

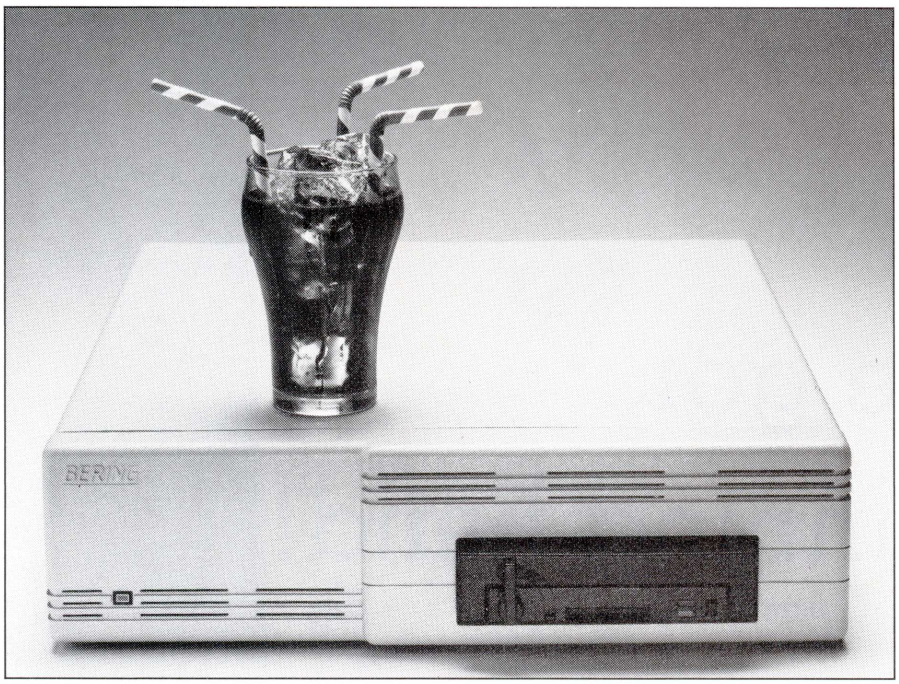

Multi-port — an HP-compatible mass storage subsystem that provides file and disk sharing capabilities for up to three users.

Ethernet to STARLAN With 487-ES Bridge

CrossComm Corporation announces a new Ethernet to STARLAN Learning Bridge. The Model 487-ES is the first of a series of high-performance devices used to interconnect dissimilar networks.

Using a proprietary software algorithm, the LAN Bridge automatically reconstructs and re-times packets so that they can be transferred from one network to another.

The 487-ES uses an 80186 processor in conjunction with an ASIC to allow performance of up to 1500 packets per second which is the maximum number of packets that can be generated on a STARLAN network. It interfaces directly to each network using standard connectors and can be downloaded from a host to provide special addressing, network management functions or programmed for special requirements such as data encryption.

The 487-ES is priced at $4,900. Contact CrossComm Corporation, P.O. Box 403, West Boylstown, MA 01503; (617) 835-4226.

Enter 903 on reader card

IEM's new Optical Disk Drive and Optical Disk Cartridge.

IEM Introduces Optical Disk Drive

IEM's 5¼-inch Optical Disk Drive for HP computers uses the CS-80 data transfer protocol, and can be used with any HP computer having an HP-IB interface — including the Series 200/300/500, HP 1000 and HP 3000 machines.

The unit can record up to 800 MB of data per cartridge. The drive itself writes on and reads from the Optical Disk cartridge within preformatted sectors using a low-power laser, and features CLV (Constant Linear Velocity) recording to yield high capacities, lineo servo with slew-mode to enhance access time, write protection, shock mounting, and write-once tape (thermal film alteration).

A software package is available that allows users to archive single files or complete disks, and programmer utilities are also available.

Prices for the Optical Disk Drive start at $14,950. Contact IEM Inc., P.O. Box 8915, Fort Collins, CO 80525; (303) 223-6071.

Enter 917 on reader card

HELLO-3000 To Debut At Interex

The HELLO-3000 application manager software product is first being introduced to U.S. users at the Interex HP Users Conference in Las Vegas.

Designed to target both the system manager and the user, it offers programmers a framework to which application programs, jobs and MPE commands are assigned.

By suspending programs after execution, HELLO-3000 increases system performance. It allows the user to page through a set of menus to find the desired application; at the same time, full security is maintained along with a logfile of all events.

Contact Alante Corporation, 5201 Great America Parkway, Santa Clara, CA 95054; (408) 986-1200.

Enter 902 on reader card

Multi-port Offers Low-Cost LAN Alternative

Bering Industries has introduced the Multi-port system, an HP-compatible mass storage subsystem that provides file and disk sharing capabilities for up to three users and a built-in, 20-MB backup.

The subsystem is available with a choice of 20-, 40-, 50-, or 70-MB Winchester drives. The fixed drive can be configured for two or three users — each allotted an equal amount of disk space on the fixed disk, and access to a 20-MB removable Winchester disk drive for easy backup or storage expansion.

Users can configure a Shared Information System (SIS) to reserve a portion of the disk for file sharing. The SIS can be accessed by only one user at a time — also a feature of the removable Winchester drive. A special traffic program prevents more than one user from editing the same document simultaneously and increases user security.

The Multi-port subsystem is compatible with HP 9000 Series 200/300/500 computers, and is housed in an enclosure that meets HP specifications and comes with all configuration software on a 5¼-inch or 3 ½-inch diskette. Standard HP-IB cabling is required, but must be purchased separately. No additional hardware or software is needed. Prices range from $4790 for Model 8220 to $6990 for the Model 8270.

Contact Bering Industries, 280 Technology Circle, Scotts Valley, CA 95066; (408) 438-8779.

Enter 904 on reader card

Spread Workstation: PCs, 3000s Work Together

L&L Products Inc., has introduced the Spread Workstation, designed for use with the PC and HP 3000. Through a remote asynchronous connection using a Hayes compatible modem, the IBM PC and compatibles can be used as a terminal for the HP 3000. Other functions include transferring files between systems error free, or performing PC-related tasks using the Workstation's built-in productivity tools.

File transfers can be performed between systems in either direction, because software is resident on both systems and can be done error-free. The software checks itself for errors and retransmits until the file transfer is correct. The PC may be used as a terminal for the HP 3000 to initiate and run programs, and terminal emulation can be suspended to perform PC-based tasks — then return to the HP and continue where one left off.

The Workstation contains the following tools: PC Spread for financial modeling, Reform to tailor data or report files to suit a specific application, full-screen editor, DOS functions, and an Autotalk feature for automating a variety of procedures and tasks. The Spread Workstation costs $900 and the HOSTLINK module for the HP 3000 costs $5,000.

They are available now from L&L Products Inc., Wheeler Professional Park, P.O. Box A-57, Hanover, NH 03755; (603) 643-4503.

Enter 905 on reader card

GATEWAY/1000 Connects HP 3000, DEC VAX, IBM

Forest Computer introduces the GATEWAY/1000, which provides intervendor, bidirectional networking services between HP 3000s, DEC VAXs and IBM mainframes.

The product was developed to provide a multidirectional solution for integrating

Continued on page 86.

Adager

Reference Manual

Installation / Update

The same procedure applies whether this is your first-ever installation or just an update. Please type four MPE commands at the system console to restore and stream the bootstrap job which installs or updates Adager:

:hello Manager.Sys :file TA; dev=tape :restore *TA; Adager :stream Adager

The Adager.Pub.Sys stream-job creates the Rego account, restores the Adager files, and even does a backup copy for your convenience. It does everything automatically and gives you explicit instructions through the system console. Since this job must execute on behalf of Manager.Sys, you have two choices regarding security. Choice (1): Temporarily delete the passwords assigned to Manager.Sys and replace them as soon as the job logs on. Choice (2): TEXT the stream-file Adager.Pub.Sys with an editor, include the passwords in line 1, keep the edited file as $NEWPASS, and then :stream $OLDPASS. As soon as you log off, $OLDPASS will disappear and your passwords will be safe.

You don't need to re-install Adager when you upgrade your MPE version or when you convert from IMAGE/3000 to TurboIMAGE.

Use

:Run Adager.Pub.Rego [; Parm=n] **(Reminder:** Co-processor Adager77.Pub.Rego must exist)

Use Parm=1 for expert mode. Use Parm=4 to force the reordering of path pointers (according to DBSCHEMA's example) whenever detail search fields have changed their relative positions. Warning: This esoteric requirement may severely degrade the performance of certain functions, such as the resequencing (or sliding) of datasets. Use Parm=8 to build a StreamFile which you can then submit as a batch job at your convenience. (You may "add" parameters: For instance, "Parm=12" combines parms 4 + 8).

Adager has a new function-interface based on verbs and objects (for example, "Fix Path"). For a detailed description, please answer "?" when Adager requests "Adager command ?"

This interface fully supports familiar one-word function names. For instance, you may type "MastCap" instead of "Change Capacity", or "PathAdd" instead of "Add Paths".

Training?

You don't need training on Adager. On the contrary: Adager will help you on theoretical and practical aspects of Image databases, right at the moment when you need the specific information and the specific understanding!

In 1978, Adager introduced an OnLine tutoring technology that quickly became famous worldwide. Simply enter a question mark "?" whenever you need help. That's all!

Use Adager to make copies of your own databases. Then, try all kinds of Adager functions on these disposable copies.

For additional bed-time reading on gossip, shop talk, facts, and opinions, you may list the files in the Doc.Rego group.

Adager research & development, distribution, maintenance, and support

Adager
Apartado 248
Antigua
Guatemala

Telephone Guatemala: +502 (2) 324333
Telex Guatemala: 5460 Adager GU

Hewlett-Packard Support Product Number: 35077.

ENTER 100 ON READER CARD

ISSUES IN Relational Technology

Considerations Of Transportability, Training And Performance

Thank goodness relational database management systems (RDBMS) exist. They now live and breathe and aren't just theoretical toys. Actual production systems with real users doing data entry, processes processing, reports reporting, graphics graphing and database structures changing are up and running with success.

Those of us who have used these tools for many years are spoiled and would rather fight than switch back to early generation languages, network databases, and having to write a program every time data needs to be converted because of some modification to the database design. Chances are that most of you know something about RDBMSs but that you're using a network system with a few add ons, like restructuring tools and added index capability.

Chances are also that you've considered the use of an RDBMS, but have hesitated because of the issues that, until recently, have clouded relational systems' prospects. Let's take a look at the issues surrounding the use of relational technology, using 1987 criteria and applying it to the RDBMSs available for the HP 3000.

[BY PAUL BASS]

Wouldn't it be marvelous if the applications developed in one environment simply could be moved to another environment with no modifications, or at worst, just a few operating system-dependent changes? Fortunately, the RDBMSs of today allow commands, applications and data structures to be moved from one environment to another.

Within the HP 3000 world, transportability usually means from a 37 to a 42 to a 52 to a 70, etc. If and when a Spectrum arrives at your door, you may want to transport to it as well.

Outside the HP 3000 world, transportability is a quagmire of mainframes, PCs, minis, workstations, etc. ADA and UNIX, however, allow software developers to provide tools that are transportable across even this wide range of hardware. This doesn't mean that anything written in ADA or UNIX is transportable. As always, think ahead to your company's future needs, do your research well and compare the capabilities of those tools.

User's View Of A Full Featured RDBMS

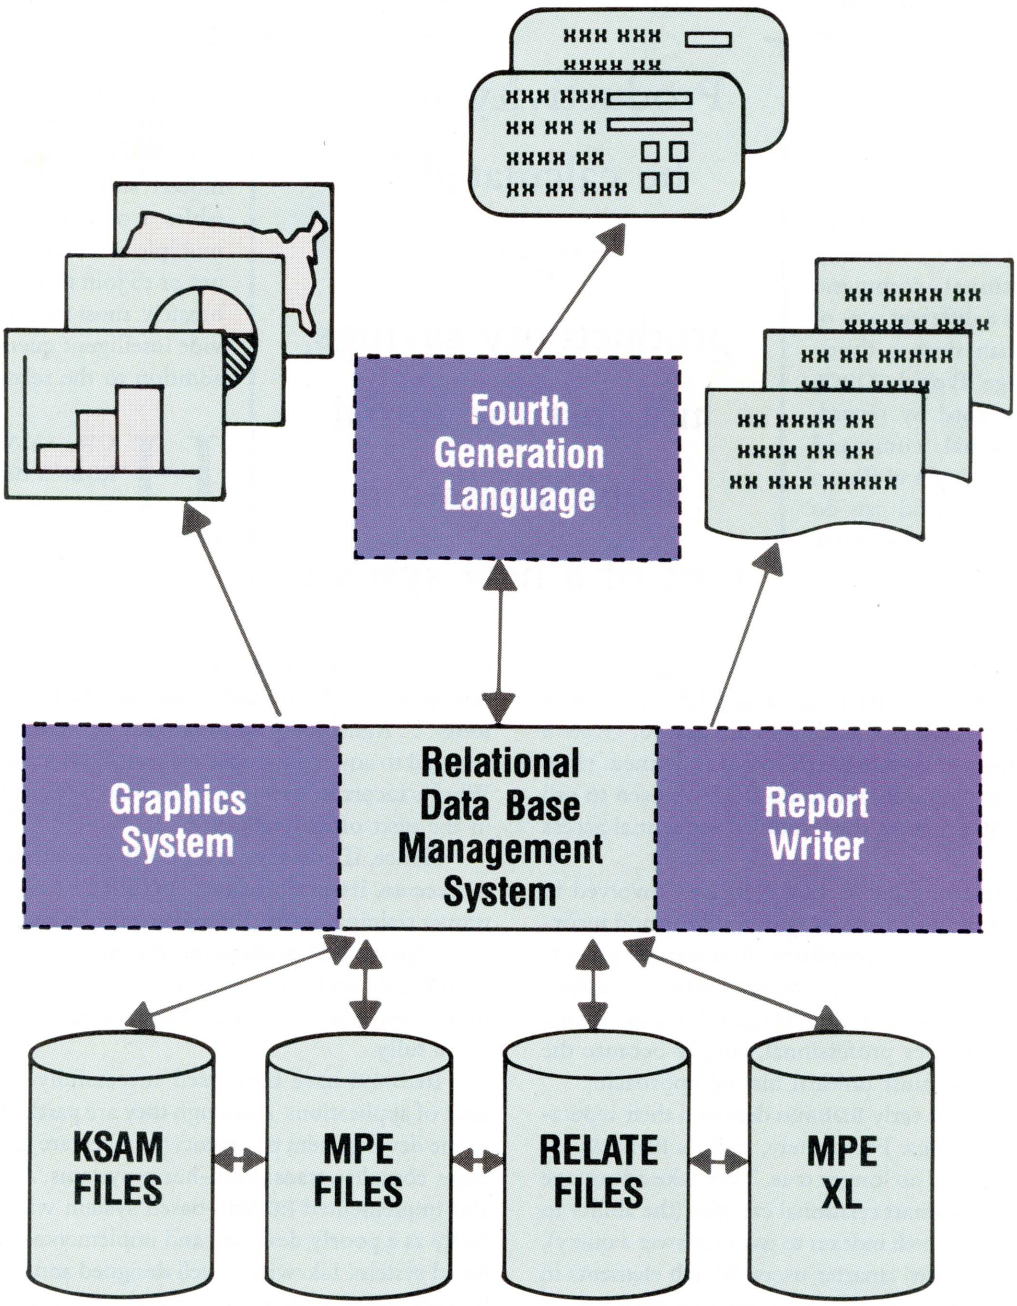

If you're currently using a network DBMS (TurboIMAGE, for example), converting to relational can be a challenge. However, the worlds of network and relational DBMSs don't have to be mutually exclusive. Some RDBMSs provide access to both database types, allowing you to use the best access method per application or to convert your systems modularly.

Similarly, the transition from third- to fourth-generation application systems can be simplified if the DBMS provides a way for a COBOL or FORTRAN application to access both existing network databases and the new relational databases.

CROSSING THE BRIDGE to the relational world requires some training because a relational command language provides a different way of accessing data than does a third-generation language. Typical COBOL programs think record by record, while the relational command language thinks in subsets of files; a single relational command can do the work of many lines of COBOL code.

For example, RELATE/3000, an RDBMS and application development system available from CRI, Inc., has a single UPDATE command that can read one file sequentially, find a match in a second file, update either or both files as defined, create a record in a third file and update the third file as defined. How much COBOL code would it take to do this? We used to call this type of COBOL program the "doubled sequential access monster."

Most of the learning curve, however, isn't involved in training, but in adapting to a new way of thinking and incorporating it into day-to-day procedures. A computer professional with many years of experience using COBOL and TurboIMAGE can have more difficulty in changing his ways than a less experienced computer professional, simply because the experienced individual must unlearn his old approach.

Unfortunately, most early RDBMSs deserved their reputations as performance hogs. Fortunately, today's RDBMSs are much smarter and this is no longer true. They take advantage of indexing, extremely smart relational calculus (the ability to decide which files and which indexes to use to answer a query), and (even more important) smarter users. Which elements in relational processing specifically can affect performance?

The first element is the parsing of the commands. Each command must be analyzed, usually character by character, to decipher exactly what you want to happen. The early RDBMSs didn't parse efficiently. However, today's RDBMSs perform parsing very quickly and cleverly.

The second is a process called relational calculus. The RDBMS must decide which file/table to read by which index and then which indexes to use to enable the joining between tables. The analysis of which indexes are available, how many records are in each file, which key provides the best access for the query, and so on, can take time. However, this up-front analysis is a small price to pay for the increased performance on each following command's execution.

> Productivity should be calculated by long-term productivity savings, not just the initial implementation cost of a new system.

Third is the execution of the commands. Without the proper selection of indexes, execution would indeed suffer. In the worst case, tables would be read sequentially multiple times to find requested values or to join to another table. Fortunately, most modern RDBMSs provide intelligent query processing in addition to the relational calculus.

HOW SMART CAN the RDBMS be? Pretty smart. In most cases, it should perform your update or respond to your query as efficiently (or more so) as any traditionally written program.

But the question is not merely the cleverness of the RDBMS. We also must ask: How smart can you be? Fortunately, we are brilliant users of our tools . . . right? Keep in mind that an RDBMS is designed to respond to any request that we give it, no matter how obscure. What is incorrect to one person could be exactly what the man in the next office wanted.

Hence, if you give it a poorly chosen set of commands to execute, its performance isn't likely to be optimal. With proper training, some cleverness and some experience, you can assist the RDBMS in choosing directions and indexes and get even better performance. Some RDBMSs even allow you to bypass the relational logic altogether to direct the operation more fully.

RDBMSs have been used successfully for all types and sizes of applications. Although they are particularly well suited to the development of prototypes, they are just as suitable for large, complex, transaction-heavy systems. A poorly designed and implemented RDBMS-based system will perform just as badly as a poorly designed and implemented network DBMS-based system. Likewise, a well designed and well implemented RDBMS based system can perform equally well while providing additional relational access and flexibility.

We all work for companies that are in business to make a profit or generate revenue for other worthy causes. The real bottom line of any computer-based system, any computer-

based tool or any DBMS is whether or not it adds to the bottom line.

I met a man recently on a business trip who was using a very old, outdated IBM system for his data processing. The hardware was so cheap that, instead of paying for maintenance, he keeps a complete machine ready to run if his machine fails. He simply throws the broken machines away and buys a replacement. He doesn't care if his machine or tools were fourth generation or second generation. I've been at other installations where thousands of dollars worth of time is wasted because a one-time investment of a few thousand dollars wasn't made to increase productivity for the application programming staff.

Productivity should be calculated by long-term productivity savings, not just the initial implementation cost of a new system. Changes, fixes, enhancements and embellishments consume as many or more resources than the initial implementation. Flexibility and control are the key issues. With RDBMSs, this is available. Reports put on the bottom of the request pile because of the effort required to produce them can be done in hours or less with their powerful query capability. Adding a forgotten field to the database takes only minutes.

There are several RDBMSs available for the HP 3000 system. RELATE/3000 is a mature RDBMS available and transportable between all HP 3000, Series 930, Apollo Workstation, Data General MV and DEC VAX computers. HPSQL/V, a relational database interface language, is available currently from Hewlett-Packard, for use on the HP 3000. ALLBASE/XL, an RDBMS with both network and relational capabilities, also from Hewlett-Packard but not compatible with HPSQL/V, will be available eventually for use on the 900 series.

The last issue is simply enjoyment. Working with a powerful, robust and flexible tool is pleasurable. I don't want to code another COBOL program again. I don't have the patience and my company doesn't have the time.

Relational technology is mature, stable and capable of high performance processing. As E.F. Codd said in the May 4, 1987 issue of *Information Week*, " . . . non-relational DBMSs have finally become obsolete." — *Paul Bass is manager of technical support at Computer Representatives Inc., Santa Clara, CA.*

Would you like to continue to see articles on this topic?
Enter On Reader Card
yes 806 no 807

A TurboIMAGE Facility To Speed Programming And Execution

The IPC Message-File Feature

[By Gilles Schipper]

Many successful HP 3000 commercial business application systems are built on the solid foundation of the TurboIMAGE database management facility available with every HP 3000 system.

Typically, the most complex programs associated with DBMS-based systems are interactive and involve several online users accessing the same database. Furthermore, a majority of these programs are designed to use IMAGE open mode 1. This provides for shared modify capability, but it's difficult to design and program and consumes the most system resources, thus executing inefficiently.

There exists an alternate design strategy involving probably the most useful file management facility introduced to the HP 3000 family of computers since IMAGE, and also one of the best-kept secrets in the HP 3000 world, namely, the Message-File feature of the Interprocess Communication (IPC) facility.

With IPC, online programs used by multiple users accessing one or more TurboIMAGE databases (as well as other files, such as KSAM) are easier to program, run faster and offer significant additional benefits, such as transparent database backup. Let's analyze the benefits associated with IMAGE/IPC design, plus associated techniques and caveats.

IPC is a subset of the MPE File System, introduced to the family of HP 3000 computers in late 1981 as part of the MPE IV operating system. IPC, though not highly publicized (a scant three pages in the *Communicator* corresponding to MPE IV), was a blessing to application system designers and programmers who wanted a simple, efficient and non-privileged method of enabling related (or non-related) processes to communicate with each other.

Before IPC, designers wanting this capability were forced to trade off ease of implementation against run-time efficiencies. IMAGE, KSAM or flat MPE files provide for relatively simple application construction and poor execution performance. Extra data segments (DS), multiple RIN (MR), and even privileged mode (PM) capabilities require very delicate and sophisticated programming abilities to gain significant execution performance advantages.

> Simply stated, IPC permits multiple user processes to communicate with each other in an easy and efficient manner. This is achieved by the use of a special kind of MPE file . . .

In most cases, however, systems were designed to include all transaction processing requirements in a single process, without considering the benefits of "distributed transaction processing" from an application design point of view. This was because the techniques available at that time created more problems than they solved. Even today, six years after the introduction of IPC, system analysts have been slow to incorporate this outstanding feature of the MPE File System into their application designs.

Interprocess Communication Facility

SIMPLY STATED, IPC PERMITS multiple user processes to communicate with each other in an easy and efficient manner. This is achieved by the use of a special kind of MPE file, known as a message file (or memory file), acting as a first-in, first-out (FIFO) queue of records. These records are placed into the file by FWRITEs and removed by FREADs.

Implementation is as simple as building the file (with the ;MSG parameter option on the :BUILD command), and programmatically issuing appropriate READs and WRITEs, as with normal MPE flat files. Additional features (such as non-destructive reads, extended waits, etc.) are accommodated easily with calls to FCONTROL. Typically, message files are associated with multiple writers and one reader, although any combination is permissible.

One very attractive feature of message files is its potential for reduced I/O overhead. If writer and reader processes are synchronized, all I/O generated by READs and WRITEs on the file actually takes place in memory (hence, the label *memory file*) instead of on disk, which is much slower.

Traditional Design

THE TYPICAL ONLINE APPLICATION on the HP 3000 today utilizes IMAGE databases in a shared modify environment. This allows multiple users to read, add, update and delete data entries in real-time to provide up-to-date information for all users of the database at any point in time.

This is the essence of modern data processing technology. The data processing state-of-the-art of a generation ago was characterized by batch processing and long delays between data entry and database updates, so that data was never really up to date.

Online shared-modify programs must open one or more IMAGE database in mode 1. Consequently, all modifications to the database must be accompanied by appropriate calls to DBLOCK. This, in turn, tends to reduce transaction throughput, because of associated overheads and delays involved and results in additional design and programming complexities due to the required thorough and delicate implementation of a proper locking strategy.

By their very nature, online, terminal I/O intensive, shared database update programs are large and complex. A design philosophy that results in the simplification of online programs and minimizes conflict resolution overhead (locking) can serve only to improve the technology of online transaction processing. Hence the marriage of IMAGE and IPC technology.

Instead of opening the database in shared modify mode (mode 1), open it with mode 6 (shared, read-only). Obviously, users no longer would be modifying the database directly in this mode. Rather, the online program would be outputting all add, update and delete transactions to a message file.

Concurrently, a batch program (no, we are not reverting to a previous generation) is responsible for performing the corresponding DBPUT, DBUPDATE and DBDELETE operations on the database. The batch process needs to open the database in mode 4 (exclusive update, allow concurrent multiple mode 6 readers).

The resulting inherent delay between transaction generation and transaction completion is volume dependent, yet in most cases, surprisingly short. Consider the advantages:

■ *Program complexity is distributed.* Rather than have your online, terminal I/O-intensive programs also responsible for database modifications and conflict coordination, this task is off-loaded to a less complex batch program, which generally is easier to debug and maintain.

■ *Overall program complexity is reduced significantly.* The online program requires no locking. The batch transaction processor program requires no locking because it has exclusive update access to the database. (There are some exceptions, described later.)

> When dealing with mode 2 opens, take care to define the database SCHEMA to ensure that security is defined at the ITEM level as well as the data set level.

■ *Execution performance is improved substantially.* Mode 4 modifications entail less overhead than mode 1 modifications because contention conflicts (normally implemented with locking) need not be handled.

■ *Terminal response times are superior.* Online users don't wait for database I/O to complete. Instead, all major I/O processing is delegated to the batch transaction processor that can be outfitted easily with a dynamic, self-tuning capability. This responds to varying overall data processing conditions, and speeds up or slows down its own execution.

■ *Database backups can be performed in a manner that's virtually transparent to the online users.*

■ *Problem resolution is simplified.* It can be far less time-consuming and, in many cases, undertaken without inconveniencing online users.

Mechanics And Techniques

DESPITE THE GENERAL OVERALL ease of implementing this design approach, it still requires some sophistication, care and ingenuity, particularly in dealing with some problems introduced with this new design philosophy. Let's identify some potential pitfalls, appropriate remedies, and some useful techniques to realize the full potential of this methodology:

1. Undesired batch process termination. In its default condition, a message file reader process will suspend (wait) upon its first message file READ, even with no writer processes (online terminal I/O programs) active. Subsequently, the writer process will suspend if there are no records available for processing and there's at least one active writer process. When all writer processes have completed (say, lunch time), the reader process will reach an end-of-file condition and presumably terminate.

To avoid this situation and eliminate the requirement to relaunch repeatedly the batch processor, two options are available. First, a call to FCONTROL with controlcode equal to 45 immediately following the FOPEN enables "extended wait" on the message file reader and prevents a physical end-of-file condition from being reached. Alternatively, though less desirable, it's possible simply to close then reopen the same file upon reaching end-of-file condition.

In either case, termination of the batch process is achieved in an orderly and predetermined manner: Create a simple FCOPY job stream to place a logical end-of-file trigger record (say, with 99 in the first two characters) that's recognized by the batch processor program.

2. Illogical transactions (such as duplicate adds, duplicate deletes, updates to records supposedly deleted, etc.). Because there's at least some delay between the generation of a transaction from the writer process and the final placement of that transaction into the database by the reader process, how do you preserve the logical integrity of the application when an attempt is made to add the same record twice? What about when an update is attempted on a record that logically doesn't exist, but physically does?

The solution is to provide a separate database to maintain transaction-in-progress records that are updated continually (by both the reader and writer processes) to reflect the logical state of the target database in real-time. This may appear to contradict the entire approach to this new design philosophy, because we're introducing another database with all the disadvantages of the traditional design, including the requirement to implement shared modify processing (that dreaded mode 1).

However, upon closer examination, this isn't the case really. In the first place, the actual database I/O associated with this transaction-in-progress maintenance requirement is very light. It requires one master data set, no chains and small-sized data entries.

Also, this technique can be used, even with the traditional design approach, to simplify database recoveries and problem resolution with minimal overhead. This additional database probably also should contain items such as tables, error messages, etc., so it can serve a larger and more general purpose. An appropriate name for this database could be SYSDB.

Finally, there exists a school of thought that discourages placing all your eggs (data) in the same basket (database). Even with the introduction of TurboIMAGE and its new (though limited) multithreading capabilities, many good arguments can be presented against defining all data requirements in a single database for significant applications. Although these arguments are beyond the scope of this article, let's consider one of them, which is germane to the issue.

By logically dividing our database into multiple physical IMAGE structures, we gain flexibility in our choice of open modes, which can directly impact overall system throughput. Simply because one or two of many data sets require online shared modify access, why should we unnecessarily impose the additional overhead requirements on the entire database? As long as we keep our database structures well-defined and logically organized, we can stand to benefit only from the divide and conquer approach to IMAGE design.

3. Real-time update requirements. There may be situations when updates to certain kinds of data must be performed in real-time. For example, consider an order entry application that allocates stock online as the order taker enters a customer order during a telephone conversation. To ensure that other orders don't deplete non-existent stock, inventories must be updated in real-time.

This problem is handled in a fashion similar to the transaction-in-progress approach. By isolating the data items requiring real-time updating in a separate database, such as SYSDB, we have our solution.

A better approach may be to have a STOKDB database that contains all warehouse inventory-related information. This database could be opened with mode 1, if new records need to be added during the day. Or mode 2 (update only, no data entry adds or deletes) could be used as long as it's understood that new inventory parts could be added only during non-prime hours, when the online portion of the application is inactive.

When dealing with mode 2 opens, take care to define the database SCHEMA to ensure that security is defined at the ITEM level as well as the data set level. Failure to do so will bring undesired results, since expected write capabilities to data items may be denied to all users except the database creator. Additional information associated with this phenomenon is left for the reader to investigate.

4. Database recovery. Due to the nature of the IMAGE/message-file design, the window of vulnerability for logical database inconsistency is slightly larger. Consequently, required database recovery, necessitated by system or application failures, becomes more complicated, though still quite manageable. A number of recovery preparation techniques can be chosen and constructed into the application itself.

Upon system or application failure, you may be required to correct the initial problem and restart the application, which continues to process the message file. However, some versions of MPE (particularly, MPE V/P delta 1 and earlier) leave open message files in a corrupted state immediately following a system failure. So, we must employ additional fail-safe mechanisms to protect the integrity of our data.

We may utilize the transaction-in-progress approach to enable us to determine exactly which transactions have been COMMITTED (output to the message file, and key information placed in SYSDB, such as order number) but not yet updated in the target database.

Once identified, we could delete (either manually with QUERY, or automatically with a user-written program) logically incomplete transactions from our database, rebuild the message file and ask the affected users to reinput the transactions, when the application is restarted.

We still have to deal with those critical inventory items that have been updated in real-time but whose corresponding transactions have disappeared into Never Never Land because of message file corruption.

An elegant and practical solution to this problem is available. Take, for example, the order entry application mentioned earlier. In the STOKDB database, we maintain, in the same data entry containing the data items stock-on-hand and qty-allocated, a data item called pending-qty-allocated.

The online program increments this field in real-time at the same time that qty-allocated is incremented. The batch processor program decrements the pending-qty-allocated field while processing the associated transaction from the message file. So, while this field is non-zero, the database is in a logically inconsistent state.

Another step in our recovery process requires us to run a program that serially processes this data set and zeros out the pending-qty-allocated field after decrementing the qty-allocated field by its value. This program can serve also as a problem-detection tool and can be referred to affectionately as the "bug-catcher."

5. Destructive read considerations. There may be a brief period from the time after the batch processor reads a transaction from the message file until it's ready to place the transaction in the database, when there is additional risk to data integrity, because of the destructive (by default) nature of message file READs.

For additional protection in those situations (necessary only for those installations that experience frequent system or application failures), it's possible to use FCONTROL (controlcode 47) to toggle alternately between non-destructive and destructive READs. This removes the possibility of a failure resulting from an incomplete database update with no corresponding source transaction in the message file. If your version of MPE is burdened with the message-file-corruption bug, this technique provides no immediate benefit.

6. Transaction logging. As many HP 3000 users know, transaction logging capability is built into the IMAGE database management subsystem. However, it would be useful, in certain situations, to be able to implement transaction logging for message files.

To do so, you could use the same logging mechanism employed by IMAGE, namely, the MPE user logging facility. Or you could use another attractive method provided by another subset of the Files System's IPC facility, namely, circular files.

Circular files are built with the ;CIR option of the :BUILD command. The essence of a circular file is that you never can exhaust its file limit, since records added to a full circular file are wrapped around to the beginning of the file in a manner transparent to the application (subject to the additional comment below).

Thus, a sequential read of the file will produce a chronologically ordered number of records corresponding to the file's capacity. So, to maintain, say, the last 5000 records' output to the circular file, you only need to build it with a capacity of 5000. All associated internal file maintenance is taken care of by MPE.

Now the comment. You should be aware of one very annoying bug associated with circular files, last encountered on MPE V/P delta-1. If you build a VARIABLE-LENGTH circular file, any attempt to write beyond the end-of-file limit will leave you with a corrupted circular file that no longer can be accessed. The work-around for that is to use fixed-length records. Or make sure you never reach end-of-file by clearing the file at regular intervals and specifying a large enough capacity to hold the maximum number of records to be accumulated during any of those intervals.

7. Transparent database backup. One very attractive benefit inherent with this design approach (at no extra charge) is the ability to perform database backups while users still are accessing the database online (Transparent Backup). To accomplish this, terminate the batch processor. Once the database is accessed by readers only, it can be backed up.

However, for those of you using TurboIMAGE (U-mit and beyond), there's yet another bug (those darn pests!) that will, for the time being, prevent you from doing such backups.

If a TurboIMAGE database has ever been opened in a mode other than 5, 6, 7 or 8, since the initial open, it's considered

to be open currently in a write-type mode, even though it isn't, until all accessors of the database exit from the application. From that point on, online users can re-enter their application, and database backup can begin, provided the batch processor isn't initiated.

It should be noted that, until backup is completed, transactions in progress will accumulate into the message file, until the batch processor is reintroduced. This means, for example, that orders placed after batch process interruption are temporarily unavailable for inquiry or modification.

By off-loading a great deal of the processing complexity to a batch process, we now have a much simpler method of identifying and resolving any application problems.

You may wish to include a facility in your application that's responsible for flagging a terminated batch process condition in the batch program itself (by placing an appropriate entry in SYSDB), and informing online users of this condition at convenient times (say, on screen form display switching).

Note: As of May 1987, this bug is fixed for U-mit (and U-mit deltas) with the T015 patch. Ask your systems engineer or the response center if you really need it. As of this writing, the T015 patch hasn't been incorporated into MPE version UB delta-2.

8. Batch processor self-tuning. It's possible to enable the batch processor program to adjust dynamically its own processing priority. This is based on external commands placed into the message file by, say, the application administrator, via a very simple FCOPY job stream similar to the batch process terminator job stream. Or it may be used on self-adjusting procedures that change their own priorities according to dynamic conditions.

The latter method is easily implemented, as follows. The online program time-stamps all transaction output to the message file. The batch processor, upon reading each transaction from the message file, compares the current time against the time-stamp associated with the transaction.

If the time difference is greater than a certain period of time (call it the upper threshold, of, say, five minutes), the program could issue a call to the GETPRIORITY INTRINSIC to increase its own priority, until such time as the time-lag is reduced to a predefined lower threshold (say, three minutes). Then another call to GETPRIORITY could be issued to revert to the original priority. The batch program needs to be prepared with PH capability to accomplish this.

Who needs OPT or APS? We've just built our own poor man's application monitor. On top of that, we now have our own perpetual motion machine.

9. Problem resolution. Many of us know how difficult it is to debug online programs. Plus, it requires two terminals in many cases, so that terminal I/O is directed to one terminal (in formatted mode) while the other is used to examine $STDLIST output.

By off-loading a great deal of the processing complexity to a batch process, we now have a much simpler method of identifying and resolving any application problems. We easily can incorporate DISPLAYs, traces and checkpoints into our batch program (transparent to any online user). Then we can examine its $STDLIST output in hardcopy form, rather than rely on reading cryptic error messages as they roll off a terminal screen.

In many cases, we can investigate and solve a problem without requiring online users to abandon their activities during diagnosis and remedial action. This is like changing a flat tire while still screeching down the highway to the bank before it closes.

10. Other considerations. You should be aware that the described method of database access (i.e., mode 6 open) is not yet common in the HP 3000 community. When you begin to implement applications employing these techniques, you suddenly may discover that some of your favorite software tools and general-purpose applications don't work as expected any more.

For example, HP Information Access (HPIA) accesses database information using open mode 5. Your IMAGE/IPC application could not coexist with HPIA access.

If you use other software products that are unable to coexist with your newly designed applications due to this easily correctable problem, you may wish to consider asking your software vendor to offer a solution in the next release.

The careful and intelligent use of a very powerful and flexible subset of MPE's file system, the IPC, provides significant new opportunities for raising IMAGE-based application design technologies to more advanced levels. Substantial improvements in programming productivity, execution performance, system availability and problem identification and resolution are at our immediate disposal. — *Gilles Schipper is president of G. Schipper and Associates, Thornhill, Ontario, Canada.*

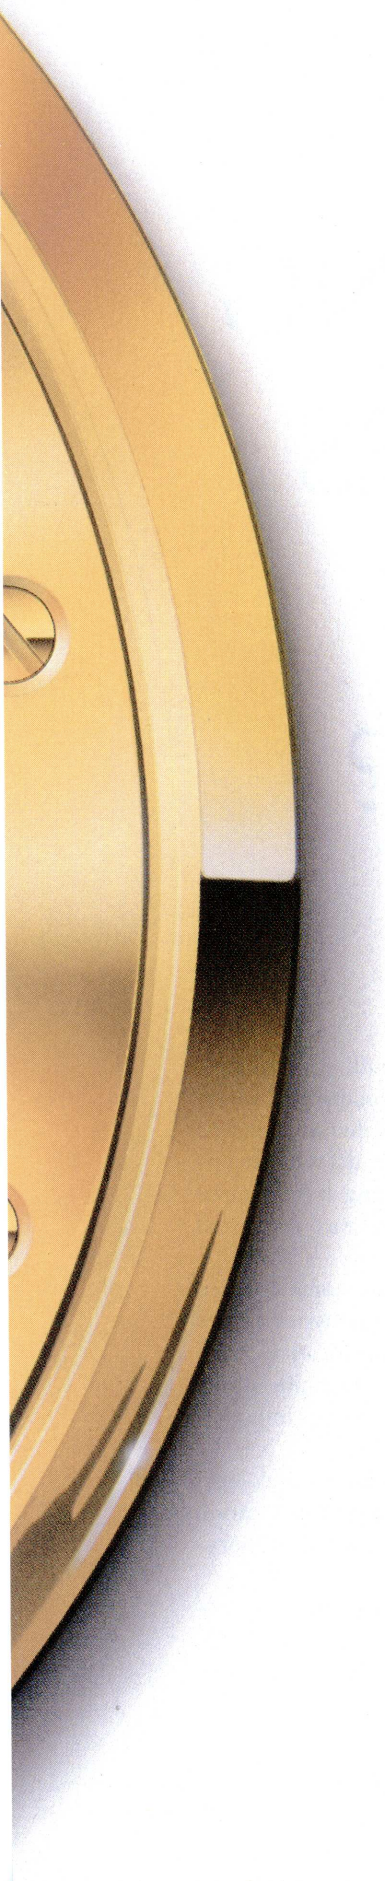

HP Networking. We connect offices, cities or countries.

Like clockwork.

An integrated business system is only as good as its connections. To other departments or offices. Or branch offices. Or even international offices.

At Hewlett-Packard, we've spent ten years designing and supporting a wide variety of local-area, wide-area and office networking solutions. All connectible to SNA-based systems. All high-performance and cost-effective. All based on OSI industry standards—so they're all compatible with other vendors. And they'll grow as you grow.

When you consider also that these solutions come from the company that never stops asking "What if...", you may wish to make a connection with us. At 1 800 367-4772, Dept. 234B.

we never stop asking "What if..."

HEWLETT PACKARD

Business Computing Systems

ENTER 226 ON READER CARD

Understanding IMAGE Internal Techniques

Do Migrating Secondaries Give You Migraines?

[BY F. ALFREDO REGO]

Secondaries are necessary evils associated with IMAGE's method of handling synonyms. *Synonyms* are necessary evils associated with hashing. *Hashing* is a technique that allows quick access to a specific entry (among millions) according to the value of its search field, regardless of the entry's location in the dataset. IMAGE's master datasets are optimized for hashed access. For online applications, serving people waiting at the counter or on the telephone, this kind of quick access provides better performance than serial scans and is simpler than indexing.

Like everything else in the universe, the advantages of hashing come coupled tightly with disadvantages. For instance, there's the possibility of synonyms: entries with different search-field values that, nevertheless, hash to the same location. Even though mathematicians have written dissertations on the desirability of perfect hashing algorithms that wouldn't produce any synonyms, every hashing algorithm known today is imperfect and produces synonyms.

A synonym chain is IMAGE's method of managing the class of entries whose search-field values hash to the same location. An entry that hashes to an empty location (or to a location temporarily occupied by a secondary) becomes a primary; an entry that hashes to a primary-occupied location becomes a secondary.

The primary has an important job: It serves as the synonym chain's ChainHead, containing information regarding the total number of synonyms in the chain, if any, and keeping track of the first and last members in the same chain. A secondary only has to worry about its predecessor, if any, and its successor, if any.

SOME SECONDARIES MIGRATE because primaries, either old or new, have top priority. Fortunately, not all secondaries migrate. The rules for migrating secondaries are clear:

■ *Whenever a new legitimate primary hashes to some location that was borrowed previously by a secondary, IMAGE migrates the secondary elsewhere, to make room for the new primary.*

■ *Whenever IMAGE deletes a primary to secondaries, it migrates the first secondary and promotes it to the primary location, to take over the ChainHead functions.*

Not all secondaries are bad. There are good ones and bad ones, according to their demands for a computer resource with direct impact on performance: disk access. Good secondaries are those that may be accessed directly in memory buffers (IMAGE's or caching's), without having to go to disk. Bad secondaries are those that force excessive disk activity.

> Like everything else in the universe, the advantages of hashing come coupled tightly with disadvantages.

Messy synonym chains, with entries scattered all over, probably will contribute numerous bad synonyms. Cleanly packed synonym chains, on the other hand, may contribute good synonyms that will be, for all practical purposes, equivalent to primary entries. Intramemory operations are faster than disk operations.

Disk caching helps, particularly if the dataset's block length is small, or if you have specified a random caching quantum larger than the default of 16 sectors. An increase in IMAGE's number of buffers (by means of DBUTIL's BufSpec command) may also help, at the expense of more memory and more CPU time to search a larger buffer pool. Under any circumstances, short synonym chains are better than long synonym chains.

THERE ARE THREE FUNDAMENTAL operations on entries: addition, deletion and finding. The object is to do these operations as quickly as possible by avoiding excessive disk access. Unfortunately, secondaries (migrating or not) tend to increase disk activity. As a consequence, the response time for online applications deteriorates and the throughput of batch jobs becomes unacceptable.

If a new entry's appointed location is vacant, simply add the new entry on the spot and mark it as occupied. The new

primary entry will have tenure for life. If a new entry's appointed location is vacant, add the new entry on the spot and mark it as occupied. The new primary entry will have tenure for life.

But if a new entry's appointed location is occupied already, you must do a lot of work. There are two possibilities:

> Smart repacking of a master dataset takes only slightly longer than an optimized serial dataset scan.

First: The current occupant *is* at its appointed place and, since it arrived before the new entry, it has seniority and tenure for life. The new entry becomes a synonym and it must be placed elsewhere as a secondary, linked to the primary by means of a synonym chain.

Before adding the entry, though, make sure that it won't duplicate an existing entry's search-field value. Therefore, you must scan the whole synonym chain before adding the new entry. If the synonym chain is long, this may take a while.

Second: The current occupant *didn't* hash to this location, but was placed here as a secondary, subject to migration. Its time to migrate has come and it must be placed elsewhere. After evicting the secondary, place the new entry in its appointed spot, as a primary.

Notice that it's easier said than done to place it elsewhere as a secondary. Finding the *elsewhere* may take a long time if the dataset has a long, uninterrupted cluster of occupied entries. This usually happens if the dataset is full or if the distribution of search-field values taxes the limits of IMAGE's hashing algorithm.

Having found the *elsewhere* doesn't guarantee tenure there, because any secondary is subject to migration in the future. If a new entry that hashes to a spot occupied by a secondary is added, the secondary must be migrated elsewhere.

If a primary with secondaries is deleted, the first secondary must be moved into the spot previously occupied by the deleted primary, because a primary is required. IMAGE, under extreme circumstances, may spend a significant amount of time chasing its own tail while migrating secondaries all over the place.

Even in a static dataset (where entries never are added or deleted), performance problems may occur when trying to find an existing entry. If the entry happens to be at its appointed location, you're in good shape. If the entry isn't at its appointed location, there are two possibilities:

- *The appointed spot is empty.* In this case, we know immediately that the entry doesn't exist, a valid result in a "find" hash.
- *The appointed location contains some other synonym entry that happens to hash to the same spot as the entry you want.* Because this entry is the primary one in a synonym chain, it keeps track of the number of synonyms. If there are no synonyms, you know that the entry doesn't exist.

If there are synonyms, scan the synonym chain until you find either the entry or exhaust chain. In either case, you may have to go to the disk more than once, depending on the length of the chain, the messiness of the chain, the dataset's blocking factor, the database's BufSpecs and disk caching.

If you're a theoretical type, you can spend endless hours figuring out the ideal mix of capacity, percent full, field-search value distribution, and so on. This is a worthy endeavor, but your dataset, most likely, is dynamic. The minute you add or delete a few thousand entries, most of your conclusions become invalid.

If you're a practical type, you might as well accept reality, bite the bullet and periodically repack your master datasets to balance the synonym load, just as you periodically repack your detail datasets, your disk drives or the drawers in your desk. If your dataset is static, you're in luck: Simply repack once.

IF THE WIDTHS OF MASTER ventilation shafts (free areas reserved for future secondaries) are too suffocating, consider changes in the dataset capacity that, together with dataset repacking, will improve the ventilation. Dataset reblocking, database BufSpec redefinition and increases in the random quantum for disk caching also may help, provided you have memory to spare.

An ideally packed master dataset allows sufficient ventilation space for DBPUT's sake (so that it doesn't take forever to find an empty spot for a new secondary or for a migrating old secondary) without having so much empty space that a serial scan (or DBSTORE) takes forever. Furthermore, the ventilation space has to be distributed intelligently throughout the entire dataset.

Fortunately, it turns out that repacking a master dataset is a very efficient operation compared to the complexities of repacking a detail dataset. Smart repacking of a master dataset takes only slightly longer than an optimized serial dataset scan.

Unfortunately, repacking doesn't seem to be an *a priori* activity. You only can repack master datasets *a posteriori*. You must repack *periodically*, and unless you quit adding or deleting entries, keep packing *ad infinitum*. — *F. Alfredo Rego, Adager, Antigua, Guatemala.*

MEGAMEMORIES™

from MARTINSOUND TECHNOLOGIES

Memory might not be very exciting, but today's software is!

Expanding the memory size of your HP Series 200/300, gives you all of the power and performance available in today's software. Martinsound Technologies' MEGAMEMORIES are the affordable way to expand the memory size of your system and still maintain the quality level that led you to an HP system in the first place.

Over 16 years ago MARTINSOUND, one of the early users of HP engineering systems, began to design and build products to meet our own need for high-technology products. We are not an importer or reseller of off-brand or off-shore clones, we use what we build... It's our efficient all domestic design, manufacturing and marketing which allow MARTINSOUND to offer the MEGAMEMORIES at prices which are as much as 60% less than comparable HP products.

A size to fit any application:

MS-310 **1 MEGABYTE**	The affordable alternative to the HP 1Mb board.
MS-320 **2 MEGABYTES**	Twice the memory, for less than the cost of an HP 1Mb board!
MS-340 **4 MEGABYTES**	The industries best Mb-per-dollar value, and you save the cost of an expansion box.
MS-380 **8 MEGABYTES**	The ultimate slot-saver! In-stock now, not just a promise for the future.
MS-350+	MEGAMEMORY upgrade for HP 350 users now available!
RAMDISK	1Mb Battery backed-up high-performance RAMDISK.

MARTINSOUND backs what it builds with a full two year warranty, 90-day money-back guarantee, and 24 hour service, if ever needed.

Whether you are a corporate, government, educational or end user, MARTINSOUND has always had the lowest price available - even in today's volatile IC market. Call us before you buy anywhere else and check it out, you owe it to yourself.

See your local Martinsound Technologies dealer, or

CALL US TODAY TO PLACE YOUR ORDER!

(818) 281-3555

"American Built Quality and Service Since 1971"

MARTINSOUND

Martinsound Technologies
1151 West Valley Boulevard, Alhambra, CA 91803-2493
Telex: 85-6573

ENTER 115 ON READER CARD

Maximizing Database Performance

Principles Of Optimized Indexing, Buffering And Data Ordering

[By Mark S. Trasko]

There's one indisputable constant in the computer world: Fast is never fast enough. Databases grow larger, while the number of users who access them and the frequency of their access continue to expand. Fast access to information through online retrievals and printed reports is crucial if database users are to get their questions answered and work completed.

Fortunately, CPU speeds on available systems are increasing rapidly. Ten MIPS and faster CPUs, once the exclusive domain of large mainframes, are becoming commonplace. Superminis and micros are providing these high speeds. For example, Sun Microsystems recently announced a 10-MIPS UNIX workstation for less than $40,000. Intel is developing the 80486 chip that reportedly will provide CPU speeds well in excess of 10 MIPS.

With CPU speeds increasing and the cost per MIPS dropping, you might think that the solution to transaction and reporting bottlenecks experienced in most database applications is around the corner. Unfortunately, it is not.

Without question, computation-intensive tasks, such as engineering, scientific and AI applications, will see dramatic increases in speed. In addition, the computation-intensive tasks required in database applications, such as sorting, will be faster. But most database applications will benefit only modestly from higher-speed CPUs because they're disk I/O-bound rather than CPU-bound. Most inquiries and reports require lengthy serial or chained database reads, even when only a small subset of records is of interest. The time required to accomplish the report depends on disk throughput rather than CPU speed. While CPU speeds continue to increase as semiconductor technology evolves, disk drives are still limited to about 30 I/Os per second, because they're mechanical devices.

Sophisticated disk controllers and operating systems can increase the aggregate I/O throughput of multiple disk systems to 100-150 I/Os per second by issuing I/O requests to several drives concurrently. Still, most database applications are I/O-bound. Even at 150 I/Os per second, today's fast CPUs are capable of executing 100,000 or more instructions for each disk access, so they spend a lot of time waiting for the next I/O to complete.

The solution is to reduce the amount of disk I/O required to get the job done. In database applications, this means that the amount of I/O needed to select and retrieve a desired subset of records for display or reporting must be cut, and the

amount of I/O required to complete a transaction must be minimized. This can be accomplished in following three ways:

- *Large block data transfers and smart buffering.*
- *Optimum data ordering.*
- *Sophisticated indexing.*

Large Block Data Transfers

A SIMPLE AND EFFECTIVE WAY to reduce I/O is to make each I/O larger, transferring more data. Although large I/Os take longer than short ones, the increase in time is slight, compared to the increase in data transferred per I/O. The maximum I/O rate of most HP disk drives is about 30 per second. That rate decreases to 15 per second when 16-KB (64 sector) transfers are used. This corresponds to a transfer rate of about ¼ MB per second. To read a file of one million 256-byte records by transferring one record per disk read requires one million/30 seconds or about nine hours. Reading the same file using 64-sector transfers takes 1,000 seconds or 17 minutes, a 30-to-one improvement.

> While CPU speeds continue to increase as semiconductor technology evolves, disk drives are still limited to about 30 I/Os per second . . .

Whenever an application needs to process all records in a file, a large block, serial read is the obvious choice. This is done by transferring data directly to the program's stack and letting the program perform the necessary deblocking of records. Or, the DBMS could do the deblocking, instead of the programmer. The DBMS usually knows what's going on in a program and can make some intelligent choices.

For example, a serial DBGET usually signals the start of a serial read, so the DBMS could use large block transfers whenever a serial DBGET is performed. If 64 sectors were read rather than the four sectors (assuming a Blockmax of 512) typically read by IMAGE, the serial read could be accomplished with 16 times fewer I/Os.

Similarly, when a chained DBGET on the primary path is done, IMAGE could read more data per I/O than usual whenever the chain length was long enough to warrant it. Assuming the data set recently had been reloaded by primary path, the longer reads would save I/O whenever several records with the same primary path key value were requested. Conversely, when directed or calculated reads are performed, a smaller I/O usually is the better choice. If only one record in a particular file area is of interest, reading multiple records wastes time and buffer space.

This concept of smart buffering extends the benefits of large block data transfers to operations other than serial reads. Disk caching, which is buffering provided by MPE, provides similar benefits, although to a lesser degree and at greater cost in performance and resources. With the exception of exhaustive serial reads of a file, all buffering techniques are imperfect. If a large block of records is read and contains only one record of interest, the extra CPU, memory and disk channel time used by the large block read is wasted. Even if a program is running on a dedicated system with no competition for memory or CPU resources, the large read and deblocking of the desired record take extra time.

Disk caching is prone to this problem. When MPE sees a read request, it has no idea whether the next read will be in the same locality. If caching is enabled, MPE does a large read into cache for better or worse. Report programs often have good disk locality and benefit from caching. Most online programs have poor locality and typically update data in the database. Since disk caching doesn't buffer write operations, only reads, online transaction-oriented programs are poor candidates for disk caching. They may run slower with disk caching enabled than with it off.

One way to mitigate the potential performance degradation of caching in online applications is to shorten the random fetch quantum. Although large fetch quantums of 64-96 sectors are best for serial reads, shorter quantums in the 16- to 32-sector range are better for most online applications, particularly with a large number of users on the system. Operational databases, such as order entry and manufacturing, tend to be highly normalized and use indexed retrieval to minimize the time required to validate transactions. As such, disk locality is low and large buffer sizes provide little benefit while costing more in memory, CPU and disk channel utilization.

GOOD DISK LOCALITY is critical to the success of buffering. Good locality means that as a process is requesting records, there's a likelihood that those records are close to each other on the disk drive. If so, one large I/O can read the first requested record and several of its neighbors. If the program then requests any of the neighboring records, they can be accessed from a memory buffer without requiring additional disk reads.

When a serial read is performed, records are read in physical order, automatically yielding good disk locality. However, getting good disk locality on keyed access takes effort. In keyed reads such as IMAGE chained reads and IMSAM sorted sequential access, records are read by key value independent of their physical order in the file. Regardless of where the next record in an IMAGE chain or IMSAM key sequence

physically resides, it's retrieved. Does this mean that the physical order of the file is unimportant?

To the contrary, the physical order of the file has a major impact on performance. If records are located physically in the same order as they're retrieved, disk locality is high and the full benefits of buffering can be attained. An IMAGE chained read of 20 records can take from one to 20 I/Os. If each record is farther away from its neighbor than the size of the buffer, each I/O will read only one of the desired records at once. Conversely, if all 20 records are contiguous physically, MPE can read all 20 into a cache buffer when the first is requested. Even if caching is disabled, IMAGE still does its own buffering. With Blockmax set to 1,024 words (my recommendation for most databases, particularly under TurboIMAGE) and a 256-byte (128-word) record size, IMAGE reads and writes eight records per disk I/O to and from internal buffers.

Because of IMAGE buffering and MPE caching, periodic reloading of an IMAGE dataset by primary path can improve the performance of reads by the primary path. Records can be ordered physically by only one path at a time, the primary path if IMAGE conventions are honored, so choosing the primary path is important. A good choice is a path by which multiple records are accessed with an average chain length of several records. A path whose chain length is usually one wouldn't be a good choice, no matter how frequently it's used.

At DISC, we've run numerous benchmarks to compare report performance on IMAGE databases with optimum data ordering versus worst-case data ordering. Worst-case data ordering is typical in operational databases unless they're reloaded periodically. Records are added as the result of transactions and thus are ordered chronologically rather than by any access path. As such, disk locality is terrible during chained access and caching can slow down a report. That's unfortunate, because caching can increase report speed dramatically when locality is high. We've seen improvements of more than five to one on reports that access a master and one or more of its details, when those details are ordered optimally.

Optimal ordering requires that details be loaded in hash sequence, the same key sequence as the master. An IMAGE-chained load/unload provides optimal ordering if the master set capacity isn't changed. Some third-party utilities also provide this function.

Keyed Sequential Access

THE GREATEST BENEFITS of data ordering result in database applications that use keyed sequential access. IMSAM and most relational databases provide keyed sequential access via an indexing structure called the B-tree. In these systems, data is accessed in sorted order for a range of key values. For example, in an accounting database, all transactions for a given time period and range of accounts might be requested. If a key is defined with account number as the primary component and transaction date as a secondary component, the desired records can be retrieved using keyed access with no serial read or sort required.

> Sorting is so important that all computer vendors and many third parties provide high-speed sorting capability.

Using direct keyed access is fast, but we can do better by recalling the previous discussion on buffering. If the data is ordered the same way as it's retrieved (in sorted key sequence), tens or hundreds of records can be retrieved per disk I/O. Furthermore, no sorting of the output is required. Reports that summarize thousands of transactions can be generated in seconds. Data can be accessed so quickly that online accounting information systems can be implemented, eliminating the need for huge printed general ledger reports. When a controller wants to review the expenditures of a department or cost center during the company's second quarter, he can enter the appropriate account and date ranges and watch as the information immediately appears on the screen.

As additional records are added to the database, they're added to the end of the file, out of sorted order. This gradually decreases disk locality and thus retrieval and report performance. To maintain optimum performance, periodic reloads can be done. The data is unloaded, sorted and reloaded, and the indexing structures are rebuilt to reflect new data ordering.

The word *reload* probably raises a red flag with nearly all DBAs, given their experience with IMAGE reloads. On large databases, reloads were time-consuming and impractical until the advent of third-party utilities such as DBMGR and Adager's DETPACK. Even with these utilities, reloads are slow, about 150,000 records per hour on a single path detail, dropping to less than 50,000 records per hour for a detail with two paths. If the superfast, high-jazz accounting information system required IMAGE-style reloads to maintain peak performance, most DBAs might decide that the users will have to wait a little longer.

Fortunately, sorting and reloading a dataset with no paths is fast. Sorting is so important that all computer vendors and many third parties provide high-speed sorting capability. HP's SORT/3000 is quite good, yielding sort speeds of 10,000 records

per minute on a Micro 3000 and twice that on a Series 70. What's really impressive is the sort speed on mainframe class machines. When we benchmarked a database on an IBM 3081 under DB2, we were amazed to see a sort speed of 150,000 records per minute using SYNCSORT. At that rate, a file of one million 128-byte records could be unloaded, sorted and reloaded in less than 15 minutes (not counting reindexing time). Even a Micro 3000 would require just two hours, practical for a night or weekend batch job.

Reindexing Speed

SINCE DATA IN LARGE FILES can be sorted and reloaded quickly, the practicality of these reloads hinges on reindexing time. KSAM files, which are B-tree structures, load at the rate of about 25,000 keys per hour for large files. At that rate, reindexing a one million record file would take 40 hours, too long to be viable. IMSAM addresses this requirement by providing a high-speed utility with a reindexing speed of one million keys per hour. This utility provides the additional benefit of allowing the DBA to select the load factor of the B-tree. Load factor is the percentage of a tree block that's occupied by keys; a high load factor conserves disk space and keeps the number of levels in the tree to a minimum. This speeds up retrievals and updates.

In summary, a million-record transaction detail with one IMSAM key can be reloaded and indexed in under four hours on a Micro 3000. Users then have instant online access to the data via account number and date. Reports that select records by ranges of account numbers and dates can be performed in seconds or minutes, rather than the hours it normally takes to generate reports from a large IMAGE database. This level of performance can be achieved easily on a compact and inexpensive Micro 3000 with one or two Eagle drives. It's cost-effective and productive to put optimized informational databases on a Micro 3000 and turn MIS users loose. They'll be in heaven, and you'll find yourself making more money.

Sophisticated Indexing

INDEXING IS ONE of the most important aspects of database design and implementation. Sophisticated indexing is the most effective way to reduce the disk I/O required to retrieve a subset of data. With advanced indexing techniques, record selections by any criteria can be accomplished using few disk reads. This permits complex selections from large databases to be done in seconds. In relational databases and IMSAM-enhanced IMAGE databases, indexing also provides a way to link files together, similar to an IMAGE path, but with more flexibility.

Consider a database containing two files. Suppose one file

contains accounting master file data and the other contains transaction detail data. If a B-tree index is created for the detail with a key consisting of account number concatenated to date, a linkage is established implicitly between the two files. This linkage is similar in functionality to an IMAGE-sorted path. Given any account number, you can use the index to retrieve all detail records for that account in date sequence. This stems from the ability to do partial key and sorted sequential access via a B-tree structure. A partial key value, the account number portion of the account-date key, is used to retrieve the first detail. Sequential reads then are used to read the remaining details for that account in date sequence.

B-tree indexes provide several additional capabilities over IMAGE-sorted paths. In this example, you're not restricted to retrieving records by just one fully specified account number at a time. You can select a range of several accounts and start and stop with any dates desired, rather than having to start at the beginning or end of a chain as with IMAGE.

Coincidentally, the update overhead of adding or deleting a key from a B-tree is the same as for an IMAGE path — four I/Os assuming a three-level tree. B-trees can be batch-loaded faster. IMSAM provides a utility that provides a reindexing speed of more than one million keys per hour. In comparison, an IMAGE master loads at the rate of less than 60,000 keys per hour because hashing distributes entries widely across the file, yielding poor disk locality. Worse yet, increasing the capacity of a master requires rehashing the entries at that slow speed. Increasing the size of a B-tree is accomplished by copying it to a larger file.

With all of these advantages, it's easy to see why B-trees are used in relational databases. Given the power, flexibility and maintainability just described, is the B-tree the state of the art in database technology? No, it's not.

THE STATE OF THE ART is better reflected by products such as OMNIDEX that include B-tree indexing capabilities and two additional features: keyword retrieval and high-speed, multifield, multiset selection. Keyword retrieval is the ability to retrieve records or documents by any combination of words and values they contain. When a field contains several words, such as a company name or product description, any of the words can be used, regardless of their position in the field.

Multifield, multiset selection extends that capability across multiple fields and sets in an IMAGE database. Selecting all customers named MARK or MARC, who work for a SYSTEMS company in DENVER, is an example of multifield selection using keyword retrieval. If a detail set containing activity data is linked to the customer master, further qualifying the above customers with selection criteria on the detail is an example of a multiset selection. All report writers can do multifield, multiset selections and typically can do string searching or pattern matching that's more general than keyword retrieval. However, such selections are slow, requiring lengthy serial and chained reads.

In contrast, these selections can be performed in seconds using its advanced index structures. In an OMNIDEX retrieval, records are qualified at the rate of 10,000 records per second per keyword (selection value), regardless of the dataset size. After selection is complete, the qualifying records can be retrieved. The following benchmark demonstrates the potential speed advantage.

> Validating a transaction requires retrievals on datasets that should be indexed heavily to permit fast and flexible access.

A test database with a 27,000-record Prospect master and a 500,000-record Activity detail was created. A report was run, selecting all USA prospects that spoke to sales rep MST in October 1986. Of the 25,000 USA prospects, only 40 met the criteria. The report took two hours on a Micro 3000 with optimum data ordering in the detail set and caching enabled and more than eight hours with worst-case data ordering.

It took up to eight hours because the report writer used the common technique of serially reading the master; then for each qualifying prospect, it did a chained read of all detail records for that prospect to see if the other selection criteria, MST and 8610, were met. Since 25,000 prospects qualified against USA and the average chain length was 20, a total of 500,000 chained reads had to be performed by the report writer. Best-case data ordering was four times faster, because of the great I/O saving. On a Series 70, with enough CPU speed to take advantage of full disk caching of optimally ordered data, the report would've been up to eight times faster.

When the report described above was run using OMNIDEX, the selection took just seven seconds and the full report was completed in 38 seconds. This represents a minimum speed improvement of nearly 200-to-1, which seems unbelievable initially, but can be explained easily. The purpose of indexing is to locate quickly a subset of records in a database so they can be retrieved. An IMAGE master record can be located with one disk read in about 35 milliseconds if

the full key value is known. Thus, indexed retrieval of a master record is about 10,000 times faster than using a serial read to find the same record. OMNIDEX is an extension of the same concept.

INDEXING ISN'T THE SOLUTION to every database performance problem. It's a cost/benefit tradeoff. It always costs disk I/Os to maintain an index when data is updated, and usually saves I/Os when data is retrieved. Improper, inefficient or incorrectly used indexing actually can slow down retrievals, or at least cost more to maintain than it's worth. The IMAGE path is a weak indexing method. It provides only full key chained access, one key at a time; no partial key, sorted sequential, keyword or multifield retrieval is provided.

Since the benefit of a path is low, a careful look at the cost/benefit equation should be made before designating more than one or two paths in a detail set. Only one path can be used at a time in a retrieval, so no multifield selection is provided by having multiple paths. Update cost is high when a search item needs to be modified because the record has to be deleted and readded, so the cost of maintaining all paths is paid even though only one was modified. It costs four I/Os per path to add or delete a record, which is about the best that any indexing structure can do. But that cost must be considered carefully when designing an application with high transaction volume. Finally, multiple path details reload slower than single path details.

Sophisticated Indexing

IN COMPARISON TO IMAGE paths, OMNIDEX uses very sophisticated indexing techniques that are low in overhead, but all indexing costs disk I/Os. OMNIDEX has the same update overhead per key or keyword as an IMAGE path, four I/Os. OMNIDEX includes a batch utility that indexes data about 10 times faster than online updates, making large database indexing very practical. A one million record dataset with eight keyword fields can be indexed in about eight hours. This is more than 10 times faster than an IMAGE detail with eight paths. The biggest difference between IMAGE and OMNIDEX, however, is on the benefit side of the equation.

What's the effect of indexing overhead on transaction volume? For maximum transaction throughput, the goal is to minimize the total I/Os required by the two components of a transaction: validation and updating. Validating a transaction requires retrievals on datasets that should be indexed heavily to permit fast and flexible access. Master files, frequently read and rarely updated, should be indexed heavily by keys and keywords to allow instant access by any words or values in any combination of fields. High transaction detail sets should be indexed lightly to minimize update overhead.

If a high transaction dataset also needs heavy indexing for reporting, there are two solutions. One is to defer indexing to a batch run during off-hours. The other is to maintain a heavily indexed informational database updated periodically from the operational database. The operational database can be optimized for online updating, while the informational database is optimized for reporting and ad hoc inquiry.

Proper application of the concepts we've explored can give users fast access to the information they need. With the advent of the Micro 3000 and Eagle drives, the option of implementing heavily indexed, highly optimized informational databases at a low cost is feasible. These databases can reside on departmental Micro 3000s, or coexist with production databases on 70-class machines. In either case, users can get fast, often instant access to data by any selection criteria without impacting production operations. The same concepts, with indexing optimized according to cost/benefit considerations, also allow maximum transaction throughput and user productivity on operational databases.

The tools and technology needed to obtain maximum performance in database applications are available. Using them increases the value of a company's data as a corporate asset, and makes all users of that data more productive. Isn't that what our jobs are all about? *—Mark S. Trasko is the president of Dynamic Information Systems Corporation (DISC) and the author of OMNIDEX.*

Getting More For Your Database Dollar

Something Old, Something New

[By Don Person]

From the user point of view, it's hard to imagine new developments in database programming, and to a certain extent, that's true. The foundation-level principles of all modern database programs were outlined years ago, and most of the rest has been filled in by advertising agencies, not programmers.

When you use application languages and prepared applications, you often use the most versatile universal designs available. By this I mean binary tree (B-tree) indices and blocked records that allow relational sort and query.

If you don't thirst for ultra-high speeds or specialized performance, off-the-shelf packages will do almost anything you want. If, on the other hand, the space-hogging tendencies and other shortcomings of fully relational database programs hamper your activities, I have some ideas for you.

Modern so-called fully relational databases as a class have a number of good points that I'll summarize:

- *Easy to set up.*
- *Simple to query, sort and define for reports.*
- *Minimal disk access required for adding or deleting records.*

Bad features:

- *Tend to be very inefficient in use of mass storage space.*
- *May be hard to recover data after random disk write failures or an index collapse.*
- *Difficult to modify after initial definition of the application.*

THE GOOD FEATURES so often outweigh the bad that we as end users tend to ignore the limitations altogether. Unfortunately, there are many types of database needs that could benefit from the relational approach, but they're very hard to code, given the constraints of modern mass storage.

While the advertising folks may have sold us on it, the rest of us don't need full relational search. Any database that uses it pays a high price in mass storage consumption for the flexibility.

Let me explain. You want to store the names of a large number of people, as well as account information. With ordinary databases, you can estimate the budget for mass storage rather simply. Add up the maximum length of each possible entry field, add a bunch more for each kind of information you store with the record, and put in a fudge factor for the indexing space for each field you define as a key.

To enable relational searching to work, the application fills unused space in each field to maintain the offset for searches. If a name and address information block adds up to a total of 400 characters, then 400 bytes of space are needed for each entry in the database, plus all the key overhead. The same amount of room is needed whether the space is filled or not.

In minicomputers and mainframes with large amounts of mass storage and tape for merge and sort, this may not be any problem at all. If you are so blessed, turn to another article. If you use a personal computer or a smaller mini-based workstation, this may interest you, so read on.

From the standpoint of pure information storage, high-powered databases can be wickedly wasteful, since on a statistical basis, most of our hypothetical record is just blank empty space. Worse is the fact that most of the relationships we're able to evaluate are of no practical use. If we're storing names and addresses, it's unlikely that we'll want to sort based on either of two large potential address fields, but the wasted space is preserved so that later fields can be traced by their offset.

In other situations, such as mail order management or inventory control, we may not care about the first name, the city, a vendor code or whatever. In other cases, we may know in advance that only one or two of many differentiated fields will ever be of interest, but we're forced to allocate space in spite of this.

Another weak spot in relational databases that use integrated record storage is the fact that a record's starting index may be:

- *Unknown or hidden by the database language, and/or*
- *Irregular, leading to trouble if recovery is ever needed.*

AWARENESS OF THESE PROBLEMS has led to a revival of interest in some rather archaic techniques. Two quaint ideas, sequential management and string indexing, are returning to minor prominence as a way of achieving higher data densities. Unfortunately, both have serious drawbacks in handling large numbers of records.

> When you use application languages and prepared applications, you often use the most versatile universal designs available.

Sequential management of keys requires that inserts and deletes move huge blocks of keys physically to perform an update. This is not bad when you know that you'll never delete or add new records, but is rather pathetic for dynamic uses where adding one new record can cause the read/rewrite of hundreds or thousands of disk sectors.

String indexing suffers from this problem and adds another of its own. Most popular operating systems place a limit on the size of a string that can be handled and thus cap the size of the database based on RAM+ program code, a big fudge factor. In my opinion, none of these revivalist ideas is worth the electrical current required to use it.

If you want efficient storage indexing for many records, let's say 20- to 100,000 or more, the only keying method that's worth looking at for hard disk use is the B-tree. That's an advantage of the relational approach, minimal disk read/write. But we don't have to go the full relational route to reap the benefits. There are ways to program the best of both worlds, if you have explicit knowledge of your needs.

Let me present an approach that offers the following:

- *Statistical compression of data fields for high density.*
- *Keyed relational access and* smart keying.
- *Inherent freedom from* record hiding *(easy recovery).*
- *Dynamic allocation with purged record space recapture.*
- *Seven Boolean flags per field free of charge (cheap anyway).*

If we store our main records in some compressed fashion, we only need keys for sorting. Thus, if our relational database only contains the key fields extracted from our record, we likely will make the fixed length component of our files much smaller than it ordinarily would be if we stored the entire record in the index. To keep track of the record, we just insert a two- or three-byte binary integer pointer into the keyed part of the records to link it. Two bytes gives us access to $2 < 16-2$ records while three bytes allows up to $2 < 24-2$ entries.

Because the keys are extracted, we have no need to fear corruption of the index; it can be rebuilt by a sequential read of the original records that are managed outside the index

> The best way to get a good fit and high storage efficiency . . . may not be blind adoption of off-the-shelf software.

anyway. As you can see, this requires access to a lower level programming language than you normally will use, perhaps, but it's the route to efficient operation.

THERE ARE MANY WAYS to compress data into less than eight bytes by selecting your range of legal characters, finding the nearest power of two greater than the highest character code you'll store and employing a bit shift algorithm to pack what you have into the least space. Additional repeated character/repeated sequence precompression would give an even higher yield, but I prefer a simpler and safer, if slightly less efficient, approach. The time to do bit compression seriously has diminishing return value anyway.

I simply remove all the trailing spaces from each field of the data and append a unique marker in between. As long as fields are prevented from containing any character greater than 128, the eight bit is a suitable flag for field termination. The bonus is that each byte used as a field separator has seven irrelevant bits left over. In my applications, I define each such free bit as a Boolean flag, and in some cases, map complete integers in the range of 0-127 into the low order part of the character. Remember that only the eight bit is significant in the overall scheme as a marker.

Here's an example using a customer file. We'll define the fields as follows:

Field#	Len.	Name
1	20	First name
2	20	Last name
3	40	Address line #1
4	40	Address line
5	20	City
6	20	State
7	10	ZIP code
8	10	Telephone number
9	60	Comments

The minimum size for each record with a typical relational database would be 240 characters. Let's look at the following entry:

```
Joe Smith
16 Outer Rd
Barnum PA 99999
501-555-1234
Sent by Vantage
```

In the compressed format I'm describing, we create a record that looks like the one here. The asterisks simply are printable substitutes for our marker tags with the most significant bit set.

```
Joe*Smith*16 Outer Rd*Barnum*PA*99999*501-555-1234*
  Sent By Vantage
```

Each byte indicated by the asterisks preserves its seven optional flags in this format:

```
Bit # : 0 1 2 3 4 5 6 7
        x x x x x x x 1
```

The **X**s denote a one or zero (true/false) condition, while the eight bit is always true and indicates that this is indeed the field delineator byte.

Here we use only 64 characters, achieving a data compression ratio of nearly four to one, with the seven marker bytes yielding 7x7 or 49 Boolean operation true/false flags. The flags commonly are used for any situation that can be described by a simple yes or no. Examples: Thinks HP is a brand of hot sauce? Has ordered products? Has blue skin? Completed high school? And so forth as your needs dictate.

Finding Flags

FINDING AN ENTIRE FLAG BLOCK is part of the routines I use to compress and restore arrays or unpacked fields. I move the flags that may have been set by the application into the marker bits and make a string from the assemblage. Then the bytes are inserted as the fields are compressed and packed together. Restoration of a record follows the reverse procedure with all the marks ending up where they started.

To find a single flag, count off the correct number of parity marked bytes until you're in range. Because this method is most useful for records with less than a full disk allocation unit, the time to pack and unpack records usually is imperceptible to real-time operators. Hint: The flags make a terrific way to indicate lock/unlock status without any SRM foolishness.

When calculating the amount of statistical compression

to shoot for, first calculate the lengths of all fields that must be constant anyway. Zip code entry data is like that. To that total, add the number of fields that the record contains. (This is the true marker overhead.) Now figure the statistical average of the sums of all other fields and add that. You've arrived now at the record length that you need.

If the value you're left with is greater than 75 to 80 percent of the relational model you would have needed, there's no gain possible. In that case, forget it! If your records have many variable length pieces of information, you can use this idea.

Two of the drawbacks to this technique should concern us. The first is the need to incorporate overflow checking in the record entry module. If we have a maximum field allowance of 400 characters and decide to use a statistical compression of 2:1, the entry module must test that the compressed length is adequate for the entry data and prompt the operator to shorten any fields that could exceed the maximum file length when summed.

In the worst case, with every field fully filled, this can add overhead to the tune of one byte per field stored. There is a fix for this problem: Instead of using a separate marker, set parity on the last character of each field. This is somewhat more hassle to *uncompact*, but guarantees that the compressed data takes no more room than the unpacked record.

The second limitation is in relational comparison for fields not defined originally as key fields when the application was created. This is not too significant, since with off-line time available, it's fairly simple to make intermediate keys based on a sequential read of the main database. The disadvantage is purely one of time, not accessibility, since the overhead we add is tied to the time needed to uncompress data for comparison. Remember, the database still has a fixed length index when the record is read, allowing an instant check.

This is just one way to get more bang for your database buck. The best way to get a good fit and high storage efficiency in your application may not be blind adoption of off-the-shelf software. Often a hybrid approach such as this can let you have your cake and eat it too. — *Don Person is an independent consultant based in Albany, NY.*

Would you like to continue to see articles on this topic?
Enter On Reader Card
Yes 800 No 801

A WINNING BET

A Guide To The North American 3000 User Conference And Las Vegas.

CONFERENCE

HP PRO Staff

Interex, the International Association of Hewlett-Packard computer users, and AZRUG, the Arizona Regional Users Group, will sponsor the Interex North American 3000 User Conference, September 20-25, in Las Vegas.

The Interex conference will be held in the Bally's-Las Vegas Hotel (formerly the MGM Grand Hotel), one of the largest luxury resort hotels in the world. The exhibit show will be open Tuesday, September 22, through Thursday, September 24.

Conference organizers expect record-breaking attendance for this show. Over 200 booths will be available to hundreds of HP users.

Bally's, 26 stories high, is located at the corner of Flamingo Road and the fabled Las Vegas Strip. It features the largest casino in Las Vegas, six restaurants, men's and women's health spas, 10 tennis courts, an Olympic-size pool and 40 stores in a shopping arcade. An 18-hole golf course is right across the street and is available to Bally guests.

The Hewlett-Packard Store, a new service for attendees, will house HP hardware and services, such as plotters, printers, PCs, tape drives, calculators, supplies, financing, etc. The Store will be held in a separate room in Bally's and will be staffed by HP factory personnel to answer technical questions about the models being shown.

We've provided a list of exhibitors and their booth numbers here, as well as a show floor plan, for your convenience. Please stop by the Professional Press booth (#501) and say hello.

During your free time, we hope you'll use the following guide to tour the sites in and around Las Vegas. The city is known for its glittering nightclubs and gambling casinos, but there also is plenty to do beyond the city limits.

See you there!

Lake Mead National Recreation Area

This 2,338-square-mile tract extends along the Colorado River 182 miles from the Grand Canyon National Park to a point below Davis Dam. It's a land of high plateaus, narrow canyons and two beautiful lakes. Impounded by Hoover Dam, Lake Mead is by volume one of the largest man-made reservoirs. It is 115 miles long, has a shore line of 550 miles at low water and 822 miles when full, with a maximum depth of 589 feet. Lake Mohave, formed behind Davis Dam, is 87 miles long. The dams are part of an irrigation, reclamation and power project of the federal government.

Hoover Dam tours are available daily. The 35-40-minute tour takes you to the dam and powerhouse. An exhibit building houses a model of the generating unit and a topographical model of the Colorado River Basin.

Self-guided tours of **Davis Dam** also are available.

Lake Mead Yacht Tours leave from the Lake Mead Marina, about seven miles east of Boulder City on NV 41 at Boulder Harbor. The tours cruise daily to Hoover Dam (about a 1¼-hour tour) on the "Echo." For fees and a tour schedule, call (702) 736-6180.

Liberace Museum

Located just over two miles east of the Strip at 1775 E. Tropicana Avenue [(702) 798-5595], the Liberace Museum displays memorabilia of Liberace's career including antique and custom cars, miniature and full-size pianos and a portion of his flashy million-dollar wardrobe.

Bally's-Las Vegas: Site of this year's Interex North American 3000 User Conference

Entrance to the exhibit costs $3.50 for adults; $3 for senior citizens; $2.50 for students; $2 for children age 6-12 and free for children under six.

Ripley's "Believe It Or Not"

Located in the Four Queens Hotel [(702) 385-4011], this museum houses more than 4,000 artifacts documenting some of the most mysterious, strangest and most ingenious phenomena imaginable. There are 10 different theme rooms and more than 1,000 individual displays.

Red Rock Canyon Recreation Lands

Pack a picnic lunch and take a drive about 15 miles west of Vegas on W. Charleston Blvd. to the Red Rock Canyon area. The view of the area's steep canyons and red and white hues of the Chinle formation and Aztec sandstone is breathtaking. There are hiking trails and nature walks led by the Bureau of Land Management ranger-naturalists [(702) 363-1921].

Spring Mountain Ranch, just three miles further west on W. Charleston Blvd., has self-guided tours, guided interpretive trails and facilities [(702) 875-4141].

Bonnie Spring Old Nevada

While in Nevada, you should try to visit at least one western mining town. This one is about 20 miles west of Vegas via W. Charleston Blvd. This historic town features a narrow-gauge railroad, museums, shops, entertainment, riding stables and a petting zoo.

Old Vegas

Located in Henderson, the fastest growing and fifth largest city in Nevada [2440 S. Boulder Hwy., (702) 564-1311], this Western theme park features shootouts, specialty rides and food. It is a replica of the Mormon Fort that was built originally in the 19th century.

Ghost Towns

The Wild West is also known for its ghost towns. Here are a few in the Las Vegas area:

Goodsprings — 30 miles southwest via I-15 to the Jean-Goodsprings turnoff (W), then seven miles west.

Sandy Valley Area — 11 miles west of Goodsprings.

Rhyolita — 112 miles northwest on US 95 to Beatty, four miles southwest on NV 58, then follow the signs.

Johnnie Mine — 67 miles northwest via US 95, 11 miles south on NV 16, then three miles east.

Wet 'n Wild

This aquatic amusement park (2600 Las Vegas Blvd. S.) features a 75-ft. water slide, rafting, flumes, water cannons, whirlpools, white water rapids, waterfalls, pools, lagoon and more. Season closes September 27.

Scenic Airlines

This is something many of us have always wanted to do: Take a sightseeing flight through the Grand Canyon. Take your camera because the 19-seat plane has wide-view, non-tinted windows for undistorted color photography.

There are four tour packages available daily. Call (800) 634-6801 or (702) 739-1900 for more information.

Interex Conference Exhibitors And Booth Numbers

ALDON Computer Group 514
Account-A-Call Corporation 605, 607
Adager 533, 535, 537, 539, 632, 634, 636, 638
Associated Systems Marketing Corporation 512
BDT Products, Inc. 127
Bering Industries, Inc. 622
Bradford Business Systems 336, 338
Bradmark Computer Systems 224, 226
Brant Computer Services Ltd. 215
Business Systems International 122
CCSC, Inc. 212
CRI, Inc. 410, 412
Cape Data 302
Carolian Systems International 133, 135, 232, 234
Chestnut Data Systems 306
Cognos Incorporated 139, 141, 143, 145, 238, 240, 242, 244
Collier-Jackson, Inc. 427, 429, 526, 528
Com Design, Inc. 109, 208
Comar, Inc. 504
Comprehensive Systems, Inc. 515, 517, 519
Computech Systems Corporation 620
Computer Data Information Systems 610
Computer Solutions, Inc. 621
Computing Options Company 600
Comshare, Inc. 112
ConAm Corporation 237, 239
Corporate Computer Systems 227, 229, 326
Cumulus Terminals 612
DATA Group Corporation, The 211
Data Comm For Business 511
Datasoft International/Cogelog 101, 103, 200, 202
Design/3000, Inc. 637
Diamond Optimum Systems 641
Dynamic Information Systems 311, 313, 315, 317, 319, 321, 323
EMC Corporation 701, 702, 703, 704
Encore Industries 129
Forest Computer, Inc. 615
Fransen/King, Ltd. 108, 110
GFK Hamburg 143, 145, 700
GTS Computer Systems, Inc. 500
Gandalf Data, Inc. 502
Gateway Systems 626, 628
Gentry, Inc. 219
Graphicus 139, 140
Great Business Solutions, Inc. 405
Hewlett-Packard Company (See Map)
High Line Data Systems 414, 416
Hillary Software, Inc. 413
Holland House 310, 312, 314
HyPoint Technology Inc. 301, 400
INTERBUSCON 235

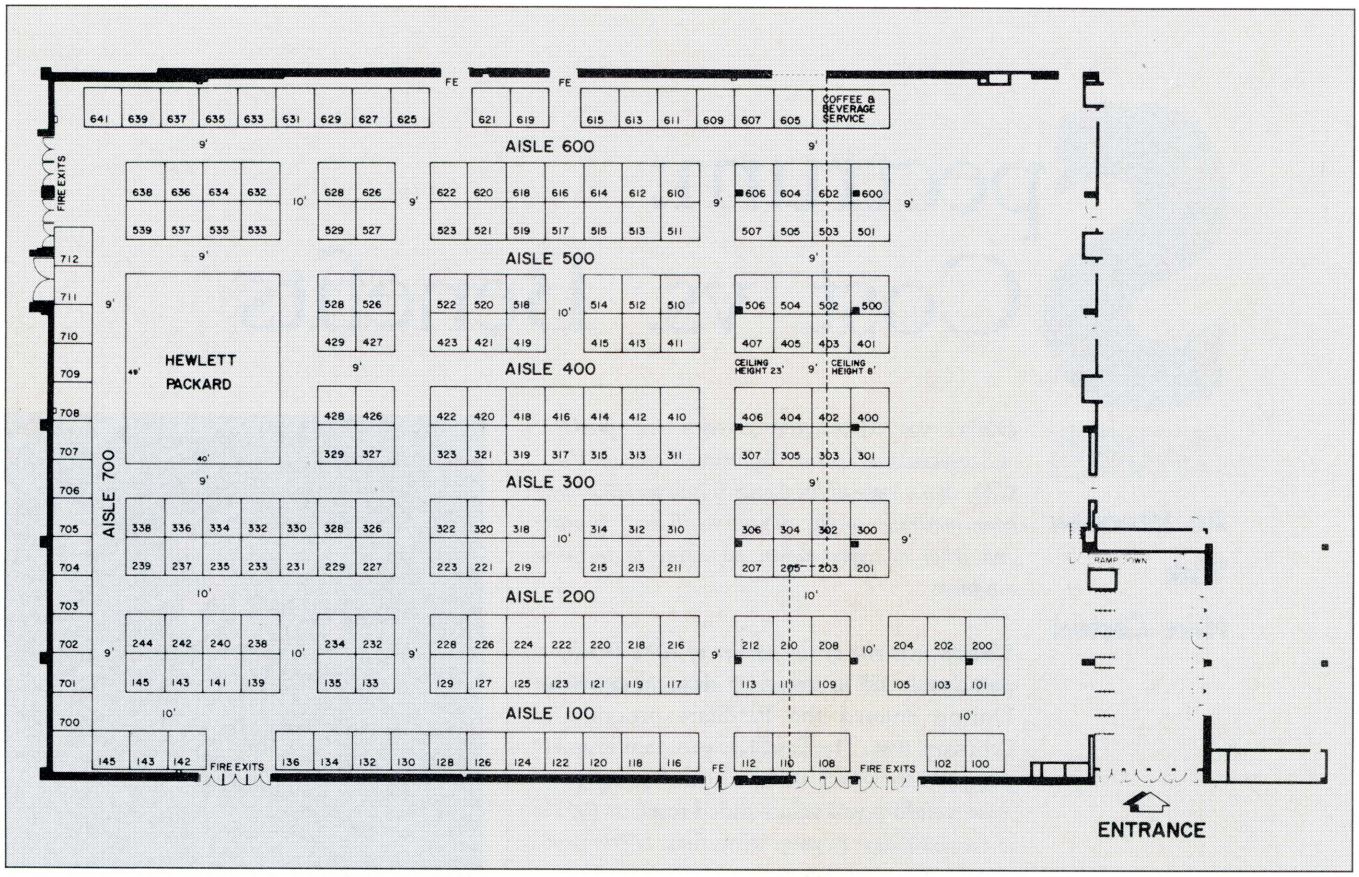

Illinois Computer Cable, Inc. 142
Imacs Systems Corporation 327, 329, 426, 428
Imperial Technology, Inc. 625
Info Centre Ltd. 332, 334, 706, 707, 708, 709, 710, 711, 712
Info Craft, Inc. 221
Infoshare Technology, Inc. 120
Innovative Software Solutions 616
John Mc Lean & Associates 222
Kelly Computer Systems 123, 125
M.B. Foster Associates Limited 418
MCBA, Inc. 318, 320
MIGHTY KEYS 627
Microtek 407, 506
MiniSoft Office Systems, Inc. 705
Mitchell Humphrey & Co. 415
NSD, Inc. 629, 631
Nedrad Corporation, The 602
O'PIN Systems 126
OPT 105, 204
Omniware International, Inc. 613
Operations Control Systems 411, 510
Peed Corporation-The Processor 619
Performance Software Group 100
Personalized Software/The Portable Paper 401
Pertaine Systems 513
Perwill Group 128, 130
Proactive Systems 132
Productive Software Systems 231, 233
Professional Press 501, 503
Protos Software Company 117, 119, 121, 216, 218, 220
RAET Software Products 639
Robelle Consulting Ltd. 322
Running Mate 210
S.T.R. Software Company 124
Second Source Computer Group 635
Smith Dennis & Gaylord 223
Softsmith 102
Software Research Northwest 116, 118
Software Systems Technology 413
Strategic Systems 614
Statistical Software Group 211, 213
Sydes N.V. 403
System Industries 203, 205, 207
Systems Express 111, 113
TELENOMICS, Inc. 618
THORN EMI Computer Software 213
Telamon 328, 330
TelemMar, Inc. 304
Teleray 134, 136
Tymlabs Corporation 419, 421, 423, 518, 520, 522
UNICOM Corporation 527, 529, 626, 628
Unison Software Inc. 420, 422
Up Time 611
VRS 228
Vesoft, Inc. 523
W.D. Farlow & Associates, Inc. 235
Walker, Richer & Quinn 303, 305, 307, 402, 404, 406
Weyerhaeuser Company 300
Where Ends Meet, Inc. 633
Wilson Publications/PCI 201
Wordware 521
Zentec Corporation 505, 507, 604, 606 ■

Spectrum: Cost vs. Benefits

An Interview With Mike Casteel

Editor's Note: Mike Casteel, executive vice president and co-founder of Unison Software (Mountain View, CA), shares some of his thoughts here on HP's Spectrum machines — the choices you'll need to make and what to expect when migrating to the new machines.

Unison was one of the earlier of HP's FastStart participants. Tell us something about the experience.
Unison entered the FastStart program in February 1986. The FastStart program is an HP program in which selected third-party software vendors and value-added resellers (VARs) get a head start in migrating their software to run on the new Spectrum class computers.

The SEMC (Software Engineering Migration Center) in Cupertino is the facility set up for testing migration, using the Spectrum machines, and getting good technical help from engineers and factory people. We have spent a total of 8-10 weeks over the past 1½ years at the SEMC, testing the development of our software accommodating migrating to native mode (NM). This required us to translate our programs from the SPL language into a NM-capable language — PASCAL, in our case.

Back in February 1986, right after our initial SEMC tour, I brought an installation tape of our TAPES product — a Tape Library Management system written in COBOL — for initial testing. We ran the program. And it worked! TAPES is a VPLUS program that accesses two IMAGE databases and a KSAM file. We ran demo databases, no problems. All VPLUS activity worked. Compatibility mode (CM) was already there. It was a pleasant surprise.

There are many MIS directors and managers out there about to make buying decisions regarding the purchase of the Spectrum Series machines. In their place, what circumstances would prompt you to buy a 900 Series machine versus a regular HP 3000?
I'd look at it as a question of obsolescence. The benefit of buying a 900 Series computer instead of sticking strictly with the "classic" 3000 is that you're looking at a longer lifetime for that machine. The "classic" 3000s are going to become obsolete via replacement with the new architecture machines over the next several years. Even if you buy a "classic" now, you're going to want to upgrade to a 900 Series anyway.

One of the things that's going to get out of date is the operating system. MPE-V on the classic machines does not have the potential and probably will not receive the kind of at-

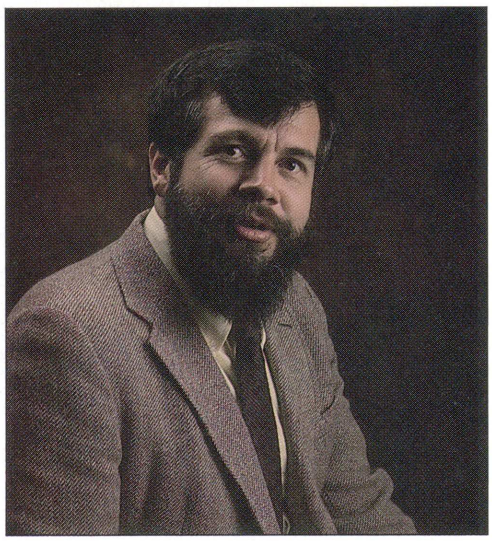

Michael Casteel.

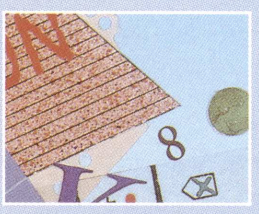

SPECTRUM

HP PRO Interview

tention in future development from HP to stretch out very many more years. MPE XL is the version of MPE that will serve for decades to come. MPE V is going to go obsolete when the Precision Architecture machines take over. We are already aware of the number of significant enhancements coming in MPE XL as soon as it's shipped.

There are enhancements not only in the user interface that are going to offer conveniences to the user — in particular, the development programmer and the DP staff — but enhancements just in the functionality of the operating system. These allow applications to be developed or extensions made to existing applications, using new operating system features that aren't there on MPE V. So, it's the new software and the new tools starting from the basic operating system and going on up that represents the usable benefit of going with the 900 Series machines versus the MPE V machines.

On the other hand, you must look at the environment you're going to be running in. Is your company going to stay in a mixed machine environment? If you have an application that needs to run on all your machines, as long as any of your machines are MPE-V based, you can't use any of the new functionalities out of MPE XL. Those functionalities are not available on MPE V; your application can't use MPE XL functionalities because then it can't run on the old MPE V machines.

But it's a different story if your applications are divided across computers. Many shops are in this kind of environment because of HP's push for distributed processing. In this case, machines are dedicated to particular applications, and a given application doesn't necessarily run on all the machines. Machine #1 runs the MRP system. Machine #2 runs the accounting system. Now if you wheel in a 950 as machine #1, all it has to run is MRP. Your MRP software can go to MPE XL functionality; it doesn't have to worry about maintaining compatibility with MPE V capabilities.

So 1) if I want to use my machine for another 10 years and 2) if I'm in an environment where I need the expandability with the new operating system functionality and 3) my application won't have to remain compatible with MPE V, then I would consider buying a Spectrum machine. I then can expand my applications and take advantage of bigger memory and the new kinds of functions — like mapped files and new databases.

In considering cost versus benefit, what kinds of issues should decision makers be concerned with regarding software migration? What sorts of questions should they be asking their manager or themselves?

As I've outlined in my Interex paper, "Spectrum Case Study: Successful Migration to MPE XL," software migration comes in several degrees of migration that can be accomplished, and there are very direct cost-versus-benefit tradeoffs associated with each.

■ *Compatibility mode.* First, there's a case of mere migration of software to execute on a Precision Architecture machine. Now, there's the big benefit offered by the compatibility mode (CM) operation of the 900 Series machines and by the operating system compatibility that HP is implementing in MPE XL. The same program object code unchanged will run on a 900 machine and an MPE V machine. That cost is very low. Cost of performing that migration is the cost of testing. The benefit is that your system will run on a 900 Series machine at virtually no cost. So, wherever you've installed a new computer, you can use the software.

Make no mistake, what we're talking about really is a migration. It's not like moving from one HP 3000 to another; you've got to go through some work.

■ *Native mode.* Second, there's migrating source code over and recompiling with the new compilers into NM. Here's where the cost benefits get interesting. The benefits are:

1. Performance of software. All code written will execute more quickly. The discount to that benefit is that most of the time your application spends executing is not your code anyway. It's executing the operating system's code to read and write files, waiting for the disk to spin around. Recompiling programs into NM doesn't offer performance benefits in that kind of environment — where most of your application performance depends on your MPE code and not your application code. It's very complicated to figure out.
2. Functionality. There are new functions made available by the operating system which are provided by NM intrinsics. These aren't directly available to old programs. If you want to use the new functionality though, you generally need to write programs and compile them with NM compilers. The new functionality isn't directly available from CM or from the old compilers.

Cost of converting programs to NM is not negligible. There is some potential for incompatibility that has to be paid attention to. The recompilation and retesting take time. And the retest is absolutely mandatory in migrating to

The "classic" 3000s are going to become obsolete via replacement with the new architecture machines over the next several years.

With OMNIDEX, Can Outperform

An astonishing claim to make. But let's look at the facts.

Fact No. 1

In data base applications, "Performance" is measured by two criteria: 1) raw throughput and 2) fast access to needed information.

Everyone knows that a 10 MIPS IBM mainframe will significantly outperform even the best of HP's hardware in transaction throughput and large reports. But most data processing departments are asked to deliver much more. As increasing amounts of corporate data are stored on-line, DP must provide specific data to a variety of users on demand. Fast access to information depends not only on available processing power, but even more on the system's ability to quickly find and retrieve the needed records from a database.

Fact No. 2

Most data base applications can be improved only slightly by increases in CPU speed.

Most inquiries and reports require lengthy serial or chained reads of the data base even when only a small subset of records is of interest. The time required to accomplish these reports depends on disc throughput not CPU speed. And while CPU speeds continue to increase as semiconductor technology evolves, disc drives, because they are mechanical devices, are still limited to about 30 I/O's per second.

To obtain fast access to information, the number of I/O's required to "get the job done" must be trimmed to an absolute minimum.

IMAGE minimizes the disc I/O needed to retrieve records by a full key value. But IMAGE presents significant limitations for most applications:

1. Masters can have only one key.
2. Details can have up to 16 keys, but look-ups can only be accomplished by one key at a time.
3. No partial key or sorted access is allowed.

So, multi-field, multi-set selections on a standard IMAGE database require lengthy serial and chained reads—imposing massive I/O demands on your system. Even an IBM mainframe, with all its processing power, is limited by the number of disc I/O's that can be accomplished in a given time frame.

There is only one way to give users instant access to the information they need: sophisticated indexing of the data.

Fact No. 3

Sophisticated indexing is the only way to dramatically decrease the disc I/O needed to retrieve information—and increase performance.

Indexing allows users to locate a desired subset of records without exhaustively scanning the data files. With advanced indexing techniques, record selections by virtually any criteria can be accomplished using relatively few disc reads. Using OMNIDEX,™ the disc I/O required for a retrieval can be reduced by a factor of 100 or more, and the time reduced to seconds. Even a powerful and expensive database like IBM's DB2,™ which uses very sophisticated techniques to minimize disc I/O, can't come close.

Fact No. 4

OMNIDEX is the most powerful DBMS available today—on any machine—at any price.

OMNIDEX provides the most sophisticated indexing capability available today, giving users incredible performance on IMAGE

Your Micro 3000 An IBM Mainframe

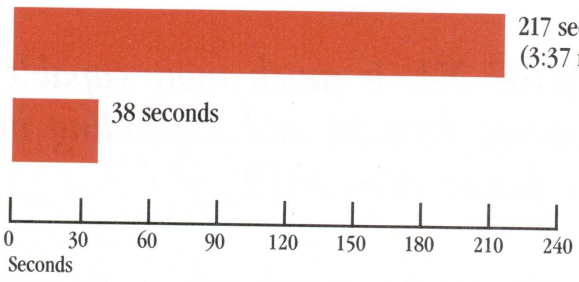

Actual benchmark study executed by an independent IBM consulting firm. A copy of the test database is available with a trial copy of OMNIDEX or OMNIQUIZ.

data bases. With OMNIDEX, users can access information in seconds by any combination of words and values across multiple fields and sets, regardless of the data base size. OMNIDEX features include:

- Record selection by multiple fields without serial reads
- Partial key retrieval and sorted sequential access
- Multiple keys in masters
- Keyword retrieval on textual data
- Easy interfacing to 3rd and 4th GL's.

For example, consider a sales support database containing a master set with 50,000 prospect records, 25,000 of which are in the USA. Linked to the master is a detail set of 500,000 activity records (e.g. prospect calls, etc.). A typical search might request all USA prospects that sales rep J. Young contacted in October of 1986. In a recent test, OMNIDEX took 7 seconds to qualify the appropriate records, and a total of 38 seconds to complete the report. Without OMNIDEX, the serial and chained reads required to accomplish the request would take at least an hour, even on a lightly loaded HP. The added horsepower of the IBM still can't match the retrieval speed of OMNIDEX. A recent benchmark executed on a dedicated IBM 3081 using DB2 required 3 minutes and 37 seconds to execute the same request—nearly *6 times longer* than OMNIDEX on the Micro 3000!

Most inquiries that users really need and want require these multi-field, multi-set selections to complete. OMNIDEX can deliver the information users need in seconds, whether your database has a 1,000 or a million records. An IBM mainframe can't even compete.

Fact No. 5

The profitability and competitive stature of your company depend on management and staff having fast access to the information they need.

Companies worldwide are now using OMNIDEX to meet that need. Hewlett Packard uses OMNIDEX in its Response Centers world-wide. Consolidated Capital uses it on their million record financial securities database.

In Rochester, New York, KODAK, Inc., uses both HP's and IBM mainframes for their sophisticated data processing needs. After OMNIDEX was installed on the company's HP 3000, one user asked the DP Manager, "When can we get OMNIDEX for our IBM mainframes?"

Mark S. Trasko
President
Dynamic Information Systems Corp.

DISC

For more information, or reprints of the above, call (303) 893-0335.

™DB2 is a registered trademark of IBM.

See us at Interex Booths #311 - #323.

A Message from DISC

ENTER 176 ON READER CARD

NM (vs. CM) because of the potential risks that something obscure has changed that impacts some particular program or some particular functionality in a program. It wouldn't be responsible to convert an application by recompiling it into NM and then assume it will work without testing it. You've just got to thoroughly test everything. And you'll find something that doesn't work and have to fix it. So, it's going to be noticeable cost to migration. It's going to take hours and hours, days and days, weeks and weeks and weeks just to recompile everything and then to test it.

From your extensive experience on the Spectrum machines, what can you say that might reassure folks out there?

Our experience on the 930s is that CM is so complete that the real migration issue is to NM. And that issue can be addressed absolutely on an as-needed basis. The hardware simply can be replaced on a 900 Series with the expectation that the migration is going to be very easy and the software systems are all going to run just as they did before. Migration to NM then can be done on a very piecemeal fashion where it's important or absolutely necessary to get the improved performance in native mode because it's a CPU-intensive application. Or get the improved functionality because the application is just unable to operate properly in the 16-bit address space of the classic HP 3000s.

You can break those bonds for those programs. In many shops, when you sit right down to look at all the software you need to run your shop, not very much of it absolutely has to be in NM in order to solve your immediate problems. Vast quantities of it don't need to be migrated to NM. You'll never notice that it's still in CM — it doesn't run that much of the time; it doesn't take that much of the CPU; it runs as fast as the disk spins and the disk spins however fast it spins; it doesn't matter what computer it's connected to.

People are hearing conflicting rumors about the operating system, MPE XL. What can you tell us about it?

I certainly can't talk about any implementation details of MPE XL. I can say, however, that the operating system is highly compatible with MPE V. That's a fundamental requirement for the success of CM. And it seems to have proven out pretty well. The new functionalities and the user interface are all moves in the right direction. Many of these new functionalities may never show up in MPE V — just because of its age and history. The new operating system offers a fresh start with a new approach to modularity in the operating system. It opens new possibilities for future enhancements. MPE XL is much more capable of moving forward and expanding in the future than MPE V.

MPE XL is much more capable of moving forward and expanding in the future than MPE V.

What can you tell us about the new database management system, ALLBASE?

HP has come out and publicly allowed that ALLBASE is going to be a great high-tech, state-of-the-art database management system someday. But existing applications are not going to want to deal with it now. ALLBASE is going to prove a base for the future, but it's not something you would look at as a migration tool or migration destination.

Most of the code out there is either COBOL or PASCAL. What can you say about the ease in migrating COBOL? PASCAL?

I expect COBOL applications are going to migrate very cleanly, from our own experience migrating TAPES and from the experiences we've heard about; i.e., Joe Alongi's report of migrating 600,000 lines of COBOL code. COBOL applications, by the way, tend to be very verbose. The number of lines shouldn't frighten you.

It makes sense that COBOL would be very easy. COBOL is a fairly well-developed language for procedural operations — for manipulating data and applying the logic. It's pretty well-designed to meet one of its original goals — machine manufacturer independence.

The HP migration isn't from operating system A to B. It's from MPE V to MPE XL. All COBOL code that doesn't address existing utilities, but just adds two numbers together and produces a result, offers no problem there. In migrating to MPE XL, your COBOL program that calls IMAGE still calls IMAGE and it all works. IMAGE is still there. Of course, adding two numbers together always has worked whenever you migrate COBOL. COBOL migration is very straightforward. I think it's going to be really easy for most programs.

PASCAL. Similar things can be said about PASCAL. However, I believe that where PASCAL is used, it's more likely to have been used to provide better access to deeper and more technical aspects of the operating system; i.e., using some of the MPE intrinsics that have changed a little bit in the way they need to be called (because of the changes in the machine's architecture). I believe those kinds of operations are more likely

to have been written in PASCAL than in COBOL. So, if you're looking for migrating PASCAL, it may well be more difficult. To the extent that the PASCAL code is not just vanilla application code, there may be more difficulty in migrating PASCAL. On the other hand, PASCAL's pretty well-defined, so if your standard language programs were simply written in PASCAL because it was your standard, those programs won't have problems migrating, either.

What is compatibility mode? Is it, as one consultant claims "a transient feature which shouldn't be counted on as the basis for long-range computing strategy?"
This is a semantic issue, I think.

Compatibility mode is one of the biggest values that HP's providing to the existing HP 3000 user as an element of the migration path over however many years it's going to take to move from MPE V — completely and without exception — to MPE XL. And that means, first of all, that you have to replace all of the MPE V HP 3000s with MPE XL 3000s. And then, you won't need compatibility mode anymore.

You then can rationally consider migrating everything to native mode. But because of costs involved, it may not make any sense to change all of your programs until such a time you decide you're going to replace the program anyway; i.e., you need a new version of the program. In that case, since you need to rewrite it anyway, now you can rewrite it purely in native mode.

You need to ask yourself, "How long am I going to use this application that has this particular program that runs once a year?" In this case, why incur any cost at all migrating it to native mode? Why not leave it in compatibility mode? It's a semantic issue. What's "long-range strategy?" What's "transient?" My point is that compatibility mode is the biggest benefit HP is offering to HP 3000 users. Why not use it?

Can you compare the migration from MPE V to MPE XL to anything you've experienced in IBM or DEC shops?
I do not have any experience with DEC migrations. I do have some with IBM. Those have been wholesale overhauls to systems. We needed to change our source code, recompile it on new compilers needed to change all our JCL (batch job control) commands. The MPE XL migration is nothing like that. Because of the operating system compatibility, it offers more favorable terms. For one thing, the JCL doesn't need changing at all — they're the same ones. Even in the migration to native mode, intrinsic functions of the operating system are still accessed in the same way; i.e., calls from the processor. So there are not changes that need to be made to the source code.

So, as opposed to some of the IBM migrations I've seen, we're not looking at the wholesale conversion of all the programs and all the jobs. It's more like a wholesale recompile of all the programs, testing, and doing a few tweaks here and there. And not having to change the jobs at all. They still work like they did before.

What can you say about the networking capability on the Series 900?
Our firsthand experience has been with the 930, operating in a LAN network which is the only initially supported link by which a 930 can be networked. It appears to be fully compatible with MPE V-based 3000s. A classic HP 3000 will network to a 930, a 930 will network to a 930, just like HP 3000s already network to each other. Our direct experience is in the areas of inter-CPU files and database access, inter-CPU command access where a user logs on over a network to a remote computer. Some of the more recent features introduced in HP's network capability — i.e., remote processing management — we haven't tried yet.

Would you like to continue to see articles on this topic?
Enter On Reader Card
yes 808 no 809

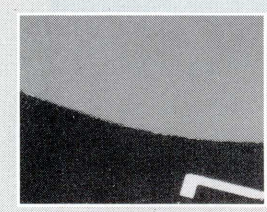

LAB REPORT

Don Person

A Tale Of The Smart And Smarter

IEEE-488 Controllers For Vectra

There are a number of vendors that boldly have moved into the IEEE-488 (HP-IB) instrument control field that HP pioneered. Each offers IEEE-488 interface support for HP Vectra and other MS-DOS computers. In addition to HP's Viper, I've chosen an add-in card representing each of the two major approaches to IEEE-488/GPIB/HP-IB control. These are of interest because they diverge from the path HP has chosen for its product.

Almost all such products have performance limitations traceable to the nastier aspects of MS-DOS and Intel 80XX performance.

Of the dozen or so competing products, most work in a similar fashion. Code is in one segment (64-KB address range) while all data references must reside in a single segment of memory. You can see that this is a potential problem for any program with variables spread over more than one 64K block. HP's Bluebird (the card offered as a disk interface; also the basis for HP's BASICA instrument library) is typical of this approach.

Another consistent aspect I found is the use of Microsoft/IBM BASICA calling conventions. While some good attempts have been made to stretch these limits, the underlying assumption seems to be that you're probably going to use these products with interpreted BASICA or so that a member of the LOTUS family or a packaged scientific, engineering or CAE program can use an HP-IB plotter.

This isn't good for those of us using compiled languages, and in particular large memory models with dynamic access of up to 640K of RAM. Here we'll concentrate on three different solutions that circumvent the small model problem.

The most expensive and most impressive of the lot is HP's Viper BASIC coprocessor/IEEE-488 controller.

Although Viper is promoted as a coprocessor, in reality it's a self-contained instrument control system with its own language. It's by far the most sophisticated.

MULTIPLE OPERATING SYSTEM computer coprocessing always has been a game of roulette, akin to entertaining a foreign guest under your own roof. In the best case, the visitor is amiable and has personal habits much like your own, sharing your living space in a state of peaceful coexistence. At worst, you could be stuck with a fanatic who makes severe demands on your resources.

The marriage of HP BASIC (courtesy of the Viper BASIC coprocessor) to the Vectra is a little like that. The Viper is no Hesbollah Shiite, but it's not exactly the perfect house guest either. After working with the card for about a month, my feelings are divided. It's hard to have no opinion about this product — most people will swear by it, or *at* it. I don 't expect much of a grey area.

If you need fast IEEE-488 computerized instrument control or you run programs already written for BASIC-speaking 200 or 300 Series HP computers and you want relatively easy file cross-access to MS-DOS, this is the card for you. You're already in the group that will swear by it.

Although Viper is promoted as a coprocessor, in reality it's a self-contained instrument control system with its own language. It's by far the most sophisticated.

As near as I can tell, Viper was designed to be an instrument controller in the classic HP mold, with none of the drawbacks of the PC-Instrument bus architecture. (PCI is promoted by HP's New Jersey Instrument division for MS-DOS/PCs.) "PC-I" and most other IEEE-488 schemes for MS-DOS can't approach the performance level achieved by Viper's high "bandwidth" parallel IEEE-488 interface.

For fast interrupt-driven control, the combination of BASIC 5.0 and its dedicated HP-IB port on Viper can't be beat. The competition challenges Viper/Vectra data transfer rates only with assembly programming. RMB BASIC 5.0 is a simple solution, and is easy to reprogram on the fly. Therein lies Viper's big advantage: IEEE-488 instru-

ment control libraries written for the 98xx Series work on it.

The way the card is put together, you can access up to about 4 MB of RAM. You'll need it. BASIC 5.0 with all the modules linked in uses nearly ¾ MB before a program is loaded. To the credit of the Corvallis Workstation Operation (CWO), a significant amount of energy has been expended showing the differences between RMB on the Viper and RMB as it runs everywhere else. All of this is appreciated, though some of it may be at the expense of explaining RMB itself. All you receive is a keyword list in the appendix of the guide. This is in keeping with the application engine thrust of the product, and I deduct no points for this. If you need to know more about RMB, you can buy plenty of extra books (listed early in the guide) to complete or augment your knowledge.

On the other hand, if you're planning to get double duty from a Vectra by doing a bit of series 200/300 BASIC 5.0 development on the side, you 'll probably end up in the swear *at* category. Double Ditto for graphics or disk intensive programs. As I proved to myself, this product is an application engine with a clear purpose, not a full-fledged computer system by itself.

A few old-line developers have cried foul over the pricing policy and the unfortunate decision to lock the Viper to Vectra alone, explicitly preventing it from running on any other AT compatible. The moaning I've heard is unwarranted, since a closer look reveals that Viper doesn't replace any other HP computer, except when used as I've indicated.

Second, by the time Viper is turned into a system, we are back up to the $6,000 price point of other 200/300 computers. Late word is that CWO will be supporting Viper on other "AT" computers, with the IBM brand probably being included by early 1988. This could be timed to coincide with the release of new software.

Viper is a fully implemented 98xx computer where the "xx"" is 16 or 36. I used 9826 files and Series 80 LIF, too, without incident. What you have is an 8 MHz 68000 system with up to about 4 MB of combined RAM/ROM possible; 512 KB RAM is the minimum. The rest fits on the little piggyback board shown in the photo. The hardware is unique, but the operating system is essentially that of a 9816/9817. Viper really has two sets of buses. The obvious one is to the Vectra/PC side, but another, the "DI/O," is perched on top of the unit. This is how direct communication is achieved with other boards in the family. Viper is envisioned as just the first in a family

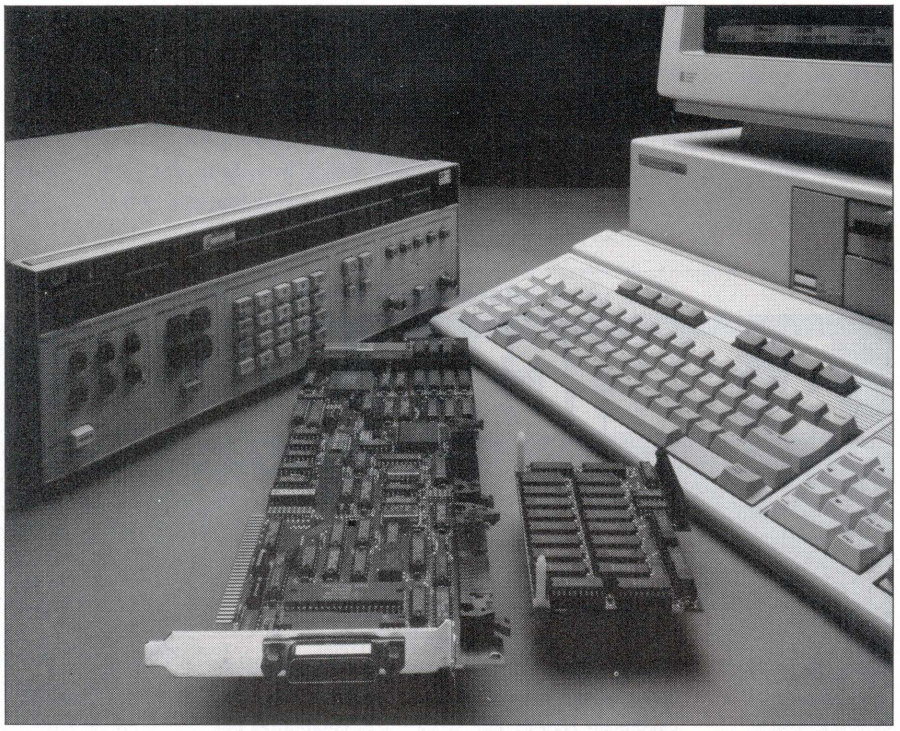

HP's Viper board and accompanying "piggyback board" for a total of 4 MB combined RAM/ROM.

engineered to fit under the hood of the Vectra/AT.

Just as vital a part of the unit is the massive interface program that coordinates I/O across the interface slot. In many ways, this is a much bigger part of the subjective "feel" of the processor than the hardware itself.

Like other 200/300 computers, the Viper can be expanded by cabling between the DI/O connectors to other modules offering SRM network support, and a true GPI/O interface to name just two. (GPI/O is part #82306A; SRM is 50963A.) Other modules are planned.

How about the combination of Viper and Vectra as a system? The cold boot process takes two minutes after you awaken old MS-DOS. If I had to dump AC power very often, I'd be unhappy. Fortunately, the warm boot time when Viper has been up and about is short. Put the boot program in your DOS AUTOEXEC.BAT file. Then when you begin work for the day, turn on Vectra, get coffee, gossip a bit and come back. You know when Viper is in control, because the screen color on the EGA and color adaptor switches from white to old familiar green.

WITH DIFFERENCES IN KEYBOARD function between the members of the 200/300 family, you'd think that this was likely to be a problem, given Vectra's version of the IBM-PC layout. Happily, all is resolved.

True, you are supplied a template to remind you of the ALT-function key sequences that replace some dedicated keys on the 200-300/98xx, but all in all, this part of the job is quite well integrated. Only minor retraining is needed. A gold star to the gang who did this part.

On mass storage the news is good, but not great. On the physical level, Series 200-300 disks follow HP Logical Interchange Format. While this defines logical file structures for different models, the common denominator until recently has been the universal adoption of a very specific disk physical format. (Not to be confused with logical layout.)

HP's older and most common physical format uses 16 sectors of 256 bytes each per track (cylinder) with an ID number of zero for the first sector on the track. Newer double-sided disks are the exception to this, as each track begins with relative sector one. This explains why double-sided 3½-inch disks return errors when you try to read them in a single-sided drive. HP file operations require access to system information which the "smart" 9121 expects to find in sector zero. (The 9122, because it doesn't use HTS addressing, can adjust to either physical format.)

The Vectra, and virtually all other MS-DOS computers using DOS 2.0 or beyond, are formatted with nine sectors per track of 512 bytes each and a first relative sector of one, and an interleave factor of one (sectors are adjacent, post-index). This leaves even mechanically compatible 5¼-inch disks looking like a totally alien medium to MS-DOS.

That's only part of the problem. The folks at CWO have done a decent job of making a 360K drive (if present in Vectra) look like an HP drive through direct reprogramming of the disk interface. A disk transfer rate test shows that it's nearly 50 percent slower loading a large file than my 9816 is working from an old reliable 82901 drive. But a tad slow is better than not at all, and I'm the first to admit that optimum interleave probably would help. Since no advice is offered on interleave control, I assume you have to experiment to find what is fast for your application. I've found default values in the HP lines of 1,2,6,7 and 8. You will probably find every other value used by somebody if you look long enough.

The 1.2-MB standard Vectra drive is a useful exchange medium in only one direction: incoming. You can't write a 5¼-inch disk for export because the

THE BOTTOM LINE

When you are looking to invest in your business for less, but want the best value for your dollar... Consider the alternative. Remanufactured equipment from Hewlett-Packard.

Peace of mind that only Hewlett-Packard can offer.

HEWLETT PACKARD

. . . Top of the Line Remanufactured Equipment from HP

See us at Interex, in the HP Showcase.

For Your Free Information Brochure, Call 800-367-4772 and Ask for Department 732P.

ENTER 242 ON READER CARD

track width is too narrow. Disk copy protection schemes don't work either. My advice to developers is to drop copy protection or risk unhappy end users.

The Viper is not suitable for running any protected 200-300 software. The disk interface to DOS is tenuous enough as it is. Copy protection schemes are more than it can handle. Two legitimate disks with copy protection simply wouldn't work on the Vectra/Viper combination, and my inference is that none ever will.

On that front, I may have found a bug in the provided LIF—> MS-DOS translator utility. When it encounters a bad sector on an HP LIF floppy disk, it may cause an incorrect write operation in a part of the ½HPW directory on the hard disk before termination. This happened to me twice.

You can store RMB files on your Vectra hard disk, but here too, there's a bit of a kludge. An odd twist is inclusion of a special MS-DOS file posing as a fake LIF style directory. This band-aid simplifies disk access for Viper, but is excess baggage to DOS. Though a file is stored as an MS-DOS byte block, it's made to appear to have an origin sector and extent (an HP LIF concept not required by MS-DOS). Files copied in by the utility routine are marked READ ONLY (no erase), a "feature" that didn't please me. This works, but it tends to obscure an unaddressed (as of this date) weakness in the system: RMB doesn't know diddly about MS-DOS hierarchical file systems (HFS). Standard MS-DOS incorporates HFS, and while it's no great shakes in my book, the details are at least public.

We are offered access to HP HFS/SRM/SDF with another add-on card, the 50963A coaxial SRM interface. HP SRM/HFS/SDF, on the other hand, is shrouded in proprietary secrecy.

The Viper RMB package needs to be taught about MS-DOS files, and must learn to treat them as an extension of its own. A logical time to make this upgrade would be with the release of HP PASCAL, slated for the end of this year or early 1988.

Miscellaneous Items:

1. The clock/timer: SET TIME worked fine, SETIMEDATE didn't work at all.
2. Programs that control the serial interface may not work correctly if they use READI/O/WRITEI/O statements. After a config for 9600, switching parameters for 300 baud in a program that worked on a 9816 didn't work on Viper. After a Reconfig and reboot, 300 baud worked, but not "on the fly." This is true if you treat the Centronics printer port on the EGA adaptor as a GPI/O or even as a true RMB printer port for other than standard output. In this case, CONTROL statements are best. I would check any program using READI/O/WRITEI/O/CONTROL statements for correctness on Viper before relying on them, no matter how well they tested on some other 98XX Series machine.

Note: I believe that only in the case of the HP-IB port is 100 percent program compatibility assured. Since this is seen by HP as the major port you're going to want, you may safely ignore the above if your application is in compliance.

NOW A TOUGH CALL. The default configuration (of four major choices) for the CRT display is probably the worst choice, though changing options for different programs can make it a little better. I wasn't pleased any time Viper had to do anything on the screen. Differences in dot resolution (Viper/9836C 512x390 vs. EGA 640,350) cause minor clipping. On the RMB side of the system, all the character attributes the CPU believes it is setting can't be handled by Vectra. This is one of several nagging compatibility questions I don't think can be done perfectly.

The difficulty seems to be that the RMB code that runs Viper thinks it's writing to its own CRT controller, when in fact it must write to outside memory

under the control of another processor. Next, there are conflicting physical limits and different read/write aperture times to contend with. Going from CRT RAM to Viper is an equal nightmare, and I suspect that this is one thing blocking adoption of a nice screen-oriented editor for RMB.

The code interface that couples Viper to Vectra's display card has got to have been difficult, and given what I know about both architectures and the gulf between them, this is still a creditable job. To summarize, it does as well as I could expect given an almost impossible task, but "pretty it is not."

Timing the ubiquitous SNOWFLAKE fractal drawing program showed Viper to be nearly 20 times slower (with an EGA display) than the 9816.

While I got two popular Series 200 third-party word processors converted and running on the Viper, the results were anything but satisfying. In one case, the program obviously uses a page redrawing algorithm after block moves that are difficult to endure. My advice? If you want word processing, buy a nice MS-DOS package and do it on the native Vectra. You'll be much happier.

Scrolling a page of alpha information, while just a flicker on the 9816, is a slow crawl up the screen. My old HP 86B can roll up almost 50 times faster than the Viper Vectra combination. That isn't good.

Hint: A DI/O-connected CRT controller module that does little more than steal power from the MS-DOS I/O connector would improve my opinion dramatically.

ROCKY MOUNTAIN BASIC is the only HP language so far for Viper, though HP PASCAL is coming. It has the same line-oriented editor found in BASIC 2.0 with additions. BASIC 5.0 retains some out-of-date concepts. OK, it's the editor! It's not the worst editor I've ever used, but it's up there with the leading contenders.

FIGURE I

Test	CEC	I/O T	Viper	Notes
File Drvr Method	960	990	11,400	Viper std ENTER/OUTPUT
CALL Method	6440	n/a	n/a	
DMA	10,300	8,800	15,700	

There's one big reason to choose RMB over other languages that run "native" on Vectra: Real-time control. Nobody else does it better than HP. Unfortunately, in other areas just about anyone but Microsoft equals or improves on it structurally. RMB has been overtaken by new language kids on the block who have improved control structures and genuine modular design.

Most important, other language vendors, less wedded to the compulsion to support antique syntax, and eager to attract new developers, have tended to be more open about access to internals.

Two Add-in IEEE-4888 HP-IB Compatible Controllers

Capitol Equipment's 4x488 is representative of designs that contain ROM routines to actually carry out IEEE-488 bus activities. Interface libraries are provided for many major languages, including True, Turbo and Quick BASIC; FORTRAN, C and even COBOL to mention a few. The interface itself and a set of alternate entry points give access to long strings, but language patch programs are left to sort out the single data segment problem I mentioned earlier.

The interface to True BASIC, for instance, copies all variables to a common segment, whether they need to be moved or not, and then copies the values back after the ROM call. While it works, it gives up a measure of efficiency, particularly with short strings. At least it does succeed in handling a problem that most of the rest of the industry hasn't come to grips with.

The IOtech GP488A is physically and electrically part of the general class of cards that package an NEC 7210 (or a similar TI 9914) IEEE-488 controller IC with a handful of extra chips to map the controller's registers into PC address space while providing bus-compatible tri-state buffering. All the smarts not found in the IC are determined by the software driver. Unlike the ROM-based system, changes in card personality do not require a change in an IC. As is the case with the 4x488, multiple DMA channel assignment is possible. When defining a system, this is imperative since there's a physical limit to the number of IRQ interrupt lines. Not every product allows the high degree of freedom that theses two offer.

The software makes or breaks these things, and here Capitol has a nice surprise. A printer redirection program lets you switch the default system printer to any of the on-board or off-board ports including the IEEE-488. Hello HP-IB printers! There's more. It allows any amount of multiple redirection including remapping a serial plotter to look like it's connected to the "488" port. It appears that any combination will work, and I tried all the obvious ones. Setting up an HP printer and redirecting PRN to it never interfered with any of the programs I ran. This is a great bonus. The ROM interface software was bug-free

MORE CACHE · LESS CASH

A Range of Add-in Memory Boards at Very Affordable Prices

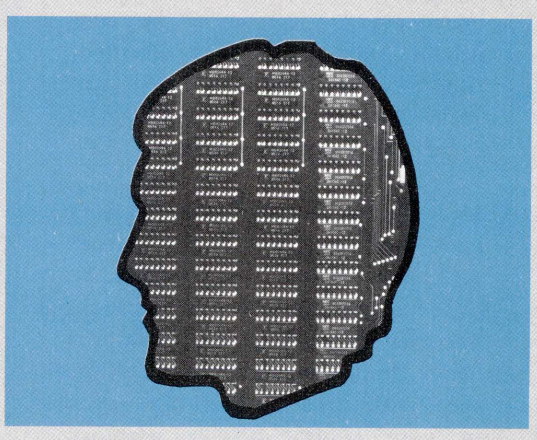

GFK

We're strong on Performance

A complete list
of things to know
about HP 3000 and HP 1000
compatible memories:

GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK GFK GFK GFK GFK
GFK GFK.

That's all.

LAS VEGAS
See us
at booth
143, 145, 700

GFK · HAMBURG · WEST-GERMANY
Ph. 40-23 34 24, Tx. 2 162 272, Fax. 40-23 17 89

GFK AMERICA, INC.
908 Providence Highway
Dedham, MA 02026/USA
Phone: 617-329 82 29
Telex: 92 0001

GFK ZÜRICH AG
Löwenstraße 17
CH-8001 Zürich/Switzerland
Phone: 1-432 61 71
Telex: 812 741 (INVE CH)

Distributors:
AUSTRIA — Phone: 222-367 53 30
BENELUX — Phone: 55-216 138
COLUMBIA — Phone: 1-217 66 00
FRANCE — Phone: 1-437 336 42
ITALY — Phone: 2-458 49 52
NORWAY — Phone: 2-13 90 90/SWEDEN
SPAIN — Phone: 1-221 82 66
UK — Phone: 905-77 99 02
VENEZUELA — Phone: 2-781 28 44

ENTER 243 ON READER CARD

when I put it through its paces with a GPIbuf.

IOtech's distinguishing feature is its software interface. The "488" hardware is quite similar to several other low-budget competing cards and by the company's admission is nothing special. A chip is a chip is a chip.

What puts IOtech closer to the head of its class is DRIVER488. It can recognize other third-party units such as CEC's older PC-488 to name but one. The installable DOS driver program that does so circumvents the "common data segment" problem I mentioned earlier. Instead of passing control by means of a CALL statement, IOtech's driver manages all but interrupt breaks through the disk operating system, making the interface look like a file to which you write commands or send data, and read command responses or receive data. It's RAM-resident and follows DOS 2.0+ conventions for file handles. Almost any program that uses files (and that hits all the ones I know of)! can interface to this. The exception would be any program that can't do DOS I/OCTL/I/OCTL$ communication. Otherwise, it's familiar, quick and very easy to program.

On the flip side, it isn't awfully fast. For most instrument work, this is no big problem. A Direct Memory Access (DMA) mode is provided so that disk and other high-speed transactions can fly along at close to the maximum interface transfer rate. Discounting DMA mode, it's the simplest to use. This counts for quite a bit, especially if you want to write short programs without becoming an HP-IB specialist. The tradeoff is some speed for the simplicity.

The tutorial content provided with IOtech and Capitol Equipment's cards is equally good, using similar analogies to group events. It's quite readable and educational.

I ran two brief tests of all three devices mentioned to see just how fast they could do their main trick; transfer data in and out on the IEEE-488 bus including some addressing overhead. The results in bps transfer are shown in *Figure 1.*

The CEC 4x488 has a limited device driver also, but it doesn't offer the complete range of function that IOtech does. IOtech doesn't have a CALL routine entry point. DMA on the 4x488 may be better owing to its true "AT" bus connection. DMA on Viper is "effective" DMA to Viper's own RAM, not Vectra's. Many Thanks to Irv Kraemer for tips on tuning the Viper DMA code.

EACH OF THE ADD-IN CARDS has very good features. Capitol Equipment's decision to offer a high-performance, multifunction card shows keen insight into what lab and engineering customers need. Though it places second in the speed trials, it's at the front of the pack in "pure" MS-DOS IEEE-488 adaptors. Packaging extended memory and extra serial/parallel I/O capability along with IEEE-488 is a clever move. That it saves slots is a bonus. It has well-documented interfaces for 640K languages, first-rate printer utilities and uses clean multilayer physical construction. It wastes neither space nor money.

On the grounds of sheer simplicity, IOtech gets the nod. If HP-IB programming isn't your passion, you'll want to add DRIVER 488 to your existing card at the very least. It flattens the learning curve, at least in programming ordinary instrument controls.

Although HP's Viper is clearly the best IEEE-488 card in my opinion, it's no more helpful in navigating the shoals of MS-DOS than any of the other systems I've looked at.

Since each system has to work with MS-DOS, one would think that producing data suitable for other DOS programs would be a priority. In the case of all but Viper, file storage is a programmer decision. The language that the card runs under makes the files, by default MS-DOS. This makes exchange easy. You have no choice but to use Viper's RMB language. It provides a complete and excellent IEEE-488 programming context, but can't handle "native" MS-DOS files.

4x488 IEEE-488,
multifunction card
Capitol Equipment
99 South Bedford St.
Burlington, MA 01803
Enter 630 on reader card.

Viper BASIC Coprocessor/
IEEE-488 controller
Hewlett-Packard
1000 N.E. Circle Blvd.
Corvallis, OR 97330
Enter 631 on reader card.

GP488A PC/IEEE-488 Card
IOtech Inc.
23400 Aurora Rd.
Cleveland, OH 44146
Enter 632 on reader card.

As a creator of HP MS-DOS file translation utilities, I'm particularly sensitive to this issue. It's outside the scope of, but implicit in the schema of the "card only" club, but is relegated to the domain of an external utility on Viper. This really should be improved ASAP.

After taking in my observations of Viper, you may think it's flawed. It's just less than perfect. Used as intended, it's an innovative solution that has a lot of potential to solve a serious problem.

HP has been, in my opinion, slow to face the fact that an explosion of third parties has eroded its position in the IEEE-488 controller business. Most of the loss has been at the hands of companies that offer accessory cards costing a small fraction of Viper. As a real-time control language, RMB has no equal, and it's here that HP now justifiably can expect to regain precious momentum in an important market. Viper offers a credible, if premium-priced, response to the burgeoning IEEE-488 card market. It comes not a second too soon for HP.
—Don Person is an independent consultant based in Albany, NY.

Would you like to continue to see articles on this topic?
Enter On Reader Card
yes 850 no 851

Copley Systems Is Hewlett-Packard's Fastest Growing Independent Dealer

. . . *Because*

Our People Make The Difference

★ Knowledgeable Sales Staff
★ Technical Support

800 Numbers staffed by software and hardware engineers

★ 2 Million Dollar Inventory
★ Unbeatable Prices

Introductory Offers

Unbeatable Low Price
Hewlett-Packard
Laser Toner Cartridges

Mfg's. List $115.00 **$75**

Calculators • Pens • Paper
Plotter Supplies • Accessories

25% OFF

Mfg's. List Prices

Immediate Delivery

CALL US TODAY!
617-449-5566

COPLEY SYSTEMS CORPORATION
185 Second Avenue, Needham, MA 02194

Outside Mass.
1-800-4-COPLEY
1-800-426-7539

ENTER 244 ON READER CARD

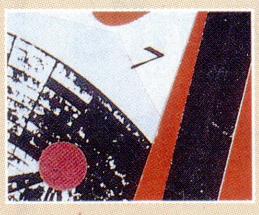

DEVELOPMENT

Lisa Burns Hartman

Don't Reinvent The Wheel

With higher emphasis on productivity and defect reduction in the MIS environment, the idea of reusing code becomes very attractive. By using standard, proven modules many times within a system, programmer productivity can be improved, since reinvention of standard functions isn't necessary. And since these reused modules already have been tested, test time is reduced and system quality is improved.

There are several methods for sharing code on the HP 3000. Here I'll compare the advantages and disadvantages of each method. Finally, by examining how one application uses these methods, I will show how they can be used effectively in an MIS environment.

By using standard, proven modules many times within a system, programmer productivity can be improved, since reinvention of standard functions isn't necessary.

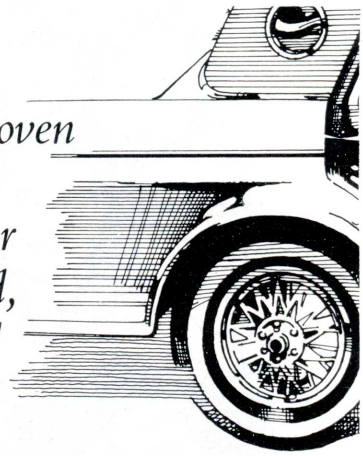

PERHAPS THE SIMPLEST APPROACH for sharing code on the HP 3000 is the use of the $INCLUDE command within the various compiler subsystems. By using the command "$INCLUDE filename", a programmer can copy the contents of the text file "filename" into a source listing at the point where the $INCLUDE command appears. This allows common file layouts and procedures to be stored in separate disk files from the rest of a program's source and shared between programs and programmers. Changes to a common file layout, for example, then can be made in only one text file instead of several affected source files. Once the change is made, the affected programs that $INCLUDE that file simply can be recompiled, and a system-wide change made very easily.

$INCLUDE files are formatted exactly the same as regular source code and can be edited with standard text editors (TDP, EDITOR, QEDIT, etc.). Since each code section that is to be copied into several programs must be stored as a separate text file, managing $INCLUDE files sometimes is tricky. However, it's far easier than attempting to make the same one-line change in 20 affected programs, and therefore is much less error-prone. Also, if you're supplying standard code or file layouts to other programming teams, you simply can give them a text file and they can use the $INCLUDE command to place it in their source listings.

A final advantage of the $INCLUDE technique is that it parallels structures available in other hardware environments. If your application runs on both HP and non-HP machines, this can be an advantage. Also, if your programming team is unfamiliar with the HP 3000, they may feel comfortable with the $INCLUDE technique because it's similar to commands they have used on other machines.

If you are programming in COBOL on the HP 3000, you can use COPYLIB files to store shared code. This technique is similar to the $INCLUDE technique, but all sections of code to be shared can be stored in one location, the COPYLIB file. Like the $INCLUDE command, COPYLIB code is copied into your source program at compile time. The command "COPY membername" is used within the COBOL compiler to copy a portion of the COPYLIB file. The "membername" parameter is stored in columns 73 through 80 of the COPYLIB file, and the compiler copies only that member requested.

The COBEDIT facility within the HP 3000 allows a group of programmers to maintain a central COPYLIB file with membernames. The file is structured as a KSAM file, making compile time access quite fast. COBEDIT is a friendly utility, but it's another text processor for programmers to learn. The ease of keeping all shared code in one location is usually worth it, however.

SUBROUTINES COMPILED as separate modules and stored in USL files allow you to isolate frequently used functions in a separate module that is PREPed into your object program. In COBOL, subprograms communicate with the calling module through a linkage area. The calling program executes the subprogram

almost like a PERFORM of another paragraph using the CALL verb: "CALL subprogram USING linkage".

Subprograms can be defined as static or dynamic. Static subprograms ($CONTROL SUBPROGRAM), once called, place data on a process' stack for the life of the process. This means that both the linkage area used to communicate with the subprogram and the subprogram's local data will remain on the stack. With dynamic subprograms ($CONTROL DYNAMIC), on the other hand, local data pops off of the stack as soon as control is returned to the calling module. Since many large applications on the HP 3000 run into stack constraints, dynamic subroutines generally are recommended in COBOL.

In other languages, subroutines stored in USL files communicate with the calling program with parameters passed by address or by value. Depending on how the parameters are passed, stack space for passed parameters may not be a problem, since variables need not be global as in COBOL. Again, subroutines in these environments may be dynamic or static. If stack space does become a problem, dynamic subroutines should be used.

Subroutines stored in USL files are a good choice for larger code routines, such as a system-wide transaction "ADD PART" or "CALCULATE BALANCE". Again, they allow programmers to make necessary code changes in only one location. Also, since they are PREPed into the object program, fewer recompiles are needed. Changes can be made by compiling only the subroutine itself, moving it into affected USL files, and rePREPing the calling programs.

Like the $INCLUDE, this technique of dealing with subroutines easily can be paralleled in other hardware and software environments and programmers generally are comfortable using it. One caution here is that programmers familiar with mainframe environments may be unfamiliar with the concept of a dynamic routine. A dynamic subroutine loses its local variables each time it returns control to the main program, however, so care must be taken to reinitialize data areas with each call and to store necessary cumulative data in global or permanent data areas for later use.

Relocatable Libraries (RLs) are simply another way of storing subroutines. Like those in USL files, subroutines stored in RLs are compiled separately and brought in to calling programs at PREP time. Several modules may be stored in an RL. They can be created as either static or dynamic, and do not necessitate recompilation of affected programs. Changes can be made by moving the new segments after they are recompiled into the RL, and then

PREPing the affected calling programs again.

All routines requested from one RL must fit in one code segment (16K words), and only one RL may be accessed during a PREP. Thus, RLs are best used for small routines. These limitations may make management of RL routines difficult. However, since they can store several routines and are accessed at PREP time, RLs can be useful for managing routines that change infrequently. By maintaining different versions of the same routine in several different RLs, RLs may be useful as a debugging aid during development.

Most HP 3000 shops I have encountered are unfamiliar with RLs and aren't comfortable using them. Since they offer few advantages over routines stored in USL files, you may wish to eliminate RL use in your environment.

Segmented Libraries (SLs) contain routines that are not linked into a program until runtime. This makes them the most flexible of all the alternatives mentioned. Routines are placed into an SL using SEGMENTER commands after the desired module is compiled into a USL file. Like RLs, SLs may contain one or more subroutines. Since they are accessed at runtime, you need only change the SL itself to effect a system change. No recompiling or rePREPing of your application programs is necessary! Subroutines in SLs are always dynamic, freeing up stack space when control is returned to the calling module.

SLs may be set up at three levels. The system SL is accessed at runtime by every program on an HP 3000 in order to call all system intrinsics. Although it is possible to place application code segments into the system SL file (SL.PUB.SYS), I would strongly advise you not to do so. When it comes time to install a new version of MPE onto your machine, a new version of the system SL file will be installed. Suddenly your user routines will disappear, and your application likely will get UNRESOLVED EXTERNAL errors at runtime.

A better plan is to set up SLs at the account and/or group levels. If your application program is located in the account MYACCT, the file SL.PUB.MYACCT also can contain SL code segments for your application. If your program is located in MYGROUP, the file SL.MYGROUP.MYACCT also can contain SL code segments for your application. The use of SLs at either the group or account level allows you to share code between programs within your own application, or with other application teams. Since you can make changes to the SL files without recompiling or rePREPing calling programs,

Program 1.

The following is an example of a COPYLIB member used in a COBOL program for a shared database access routine.

```
04707           187100 COPY DBFIND.
04708 DBFIND    001000*
04709 DBFIND    001100*  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *
04710 DBFIND    001200*
04711 DBFIND    001300*  THIS PARAGRAPH IS USED TO LOCATE A MASTER SET ENTRY
04712 DBFIND    001400*  USING THE SPECIFIED SEARCH ITEM.  IT SETS UP THE
04713 DBFIND    001500*  POINTERS FOR CHAINED ACCESS THRU THE SPECIFIED DETAIL.
04714 DBFIND    001600*
04715 DBFIND    001700*  INPUTS:   ICP-BASE
04716 DBFIND    001800*            ICP-DSET
04717 DBFIND    001900*            ICP-ITEM
04718 DBFIND    002000*            ICP-ARGUMENT
04719 DBFIND    002100*            COME-FROM-NUM
04720 DBFIND    002200*
04721 DBFIND    002300*  OUTPUTS:  NON-ERROR STATE        ERROR STATE
04722 DBFIND    002400*            ===============        ===========
04723 DBFIND    002500*            ICP-STATUS             ICP-STATUS
04724 DBFIND    002600*                                   ERROR-FLAG
04725 DBFIND    002700*                                   ERROR-MSG
04726 DBFIND    002800*                                   COME-FROM-NUM
04727 DBFIND    002900*
04728 DBFIND    003000*  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *  *
04729 DBFIND    003100*
04730 DBFIND    003200 DB-FIND.
04731 DBFIND    003300
04732 DBFIND    003310$IF XO=ON
04732 DBFIND    003320     DISPLAY "YOU'VE FOUND DBFIND".
04732 DBFIND    003330$IF
04732 DBFIND    003340
04733 DBFIND    003400     MOVE SPACES TO ERROR-FLAG.
04734 DBFIND    003500
04735 DBFIND    003600     CALL "DBFIND" USING ICP-BASE
04736 DBFIND    003700                         ICP-DSET
04737 DBFIND    003800                         ICP-MODE1
04738 DBFIND    003900                         ICP-STATUS
04739 DBFIND    004000                         ICP-ITEM
04740 DBFIND    004100                         ICP-ARGUMENT.
04741 DBFIND    004200
04742 DBFIND    004300     IF ICP-COND-CODE = 0
04743 DBFIND    004400         GO TO DB-FIND-EXIT.
04744 DBFIND    004500
04745 DBFIND    004600     IF ICP-COND-CODE = 17
04746 DBFIND    004700         MOVE "E" TO ERROR-FLAG
04747 DBFIND    004800         MOVE "NO MASTER ENTRY" TO ERROR-MSG
04748 DBFIND    004900
04749 DBFIND    005000     ELSE
04750 DBFIND    005100         MOVE "I" TO ERROR-FLAG
04751 DBFIND    005200         MOVE ICP-BASE TO ERR-FILE
04752 DBFIND    005300         MOVE "DBFIND FATAL ERROR" TO ERR-MSG
04753 DBFIND    005400         MOVE ICP-STATUS TO IMAGE-GEN-STATUS
04754 DBFIND    005500         GO TO 9999-FATAL-ERROR-EXIT.
04755 DBFIND    005600
04756 DBFIND    005700 DB-FIND-EXIT. EXIT.
```

Program 2.

The following is an example of a $INCLUDE file used in a PASCAL program to define procedures for a shared error routine.

```
0    67.000  0  $INCLUDE 'ERRMODY'$
0     1.000  0  {*************************************************}
0     2.000  0  { $INCLUDE ERRMODY - procedure heading exports    }
0     3.000  0  {                                                 }
0     4.000  0  { Owner Module:        ERRMOD - error handling module }
0     5.000  0  {                                                 }
0     6.000  0  { Description:         Procedure heading definitions }
0     7.000  0  {                      for error module.          }
0     8.000  0  {                                                 }
0     9.000  0  { Maintenance History:                            }
0    10.000  0  {                                                 }
0    11.000  0  {E/C Ref #  Author          date   Description    }
0    12.000  0  {                                                 }
0    13.000  0  {           JRM           11/19/86 Initial Release. }
0    14.000  0  {*************************************************}
0    15.000  0
0    16.000  0  PROCEDURE init_errmod(
1    17.000  0        mod_eset_num      : short
0    18.000  0    ); EXTERNAL;
0    19.000  0
0    20.000  0  PROCEDURE fatal(
1    21.000  0        error_set         : short;
2    22.000  0        error_number      : short;
3    23.000  0        error_data        : short
0    24.000  0    ); EXTERNAL;
0    25.000  0
0    26.000  0  PROCEDURE file_err_disp(
1    27.000  0        file_num          : short
0    28.000  0    ); EXTERNAL;
0    29.000  0
0    30.000  0  FUNCTION file_gen_msg(
1    31.000  0        file_error_num    : short
0    32.000  0    ): str72; EXTERNAL;
0    33.000  0
0    34.000  0  FUNCTION errmod_msg_gen(
1    35.000  0    VAR message_num       : short;
2    36.000  0    VAR message           : pac64
0    37.000  0    ) : boolean; EXTERNAL;
0    38.000  0
0    39.000  0  {*************************************************}
0    40.000  0  {                    END OF ERRMODY               }
```

Put a Falcon in your HP 3000 and watch it fly.

Introducing Falcon Series™ high speed disk subsystems from EMC.

The Falcon Series is the first family of disk drive subsystems that outperforms HP's drives in every respect.

Nothing comes close to the Falcon for speed. Because every subsystem is equipped with 4MB of high speed cache, a dedicated Motorola 68000 microprocessor and the most efficient caching algorithm possible. Together with the fastest Winchester disk drives, they can improve on the speed of an HP 7933 disk drive by a full 90%. They even beat HP's new Eagle drives.

The Falcon also lets you expand your storage without expanding your computer room. By combining 406MB or 812MB subsystem modules in a single cabinet, you can put up to 2.4 gigabytes of storage in the same footprint as an HP 7933/5 disk drive. That's 6 times the storage in the same space.

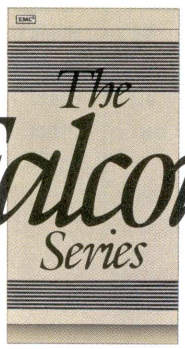

What's more, no one can give you a more reliable drive. With a MTBF of up to 30,000 hours, our sealed Winchester drive technology is 3 times as reliable as older removable drives. They give you more environmental flexibility too, since they require no special air conditioning or ventilation. And maintenance is not a problem. Your contract with HP remains unaffected. EMC offers you several maintenance plans designed to make your disk management simple and easy, and maximize your uptime.

To cap it all off, Falcon subsystems cost you less per megabyte than comparable HP units.

So if you want your HP3000 to fly, give it a Falcon.

For more information, contact EMC Corporation, Hopkinton, MA 01748-9103. Or call **1-800-222-EMC2.** (In MA, 617-435-2541.)

EMC²

The System Enhancement Company

FIGURE 1

$INCLUDE files

—Good for data definitions and for small procedures
—Separate file for each shared module or file layout
—Requires recompile of all affected programs for changes

COPYLIB files

—Available only in COBOL environment
—Central location of shared code makes version control easier
—Requires recompile of all affected programs for changes

Subroutines in separate USL files

—Good for larger procedures, frequently-used functions
—May be static or dynamic
—Dynamic modules must initialize data each time, save cumulative data
—Requires recompile of changed modules only, re-PREP of calling modules

RL files

—May be static or dynamic
—All requested routines must fit in 16K words or less
—Requires recompile of changed modules only, re-PREP of calling modules
—Good for small routines which do not change very often

SL files

—Always dynamic
—May be set up at Group or Account levels
—Accessed at run-time, no recompile or re-PREP required
—Excellent for routines that change frequently

SLs are an excellent choice for routines that are likely to change. For example, you can change a module you supply which accesses a common database and ensure that other applications also are accessing the correct version of the module.

TO ILLUSTRATE THE USE of each of these techniques, let's look at a real application in action. The project I work on currently is a large (approximately 160,000 total lines) COBOL application that uses V/PLUS calls for screen I/O and IMAGE/3000 for database access. We also have a few modules written in PASCAL. The system developers, in their wisdom, made use of all of the above techniques except RLs in order to share common code, and we continue to do so!

Since we are COBOL-based, we can take advantage of the COPYLIB facility. About 60 percent of our COPYLIB file is used for data definitions, and all of our commonly used data areas are set up as COPYLIB members. For example, we have a copy member for each of our own dataset layouts. We have copy-members for data we read from other teams' databases, like customer and product file databases. File layouts used for interfaces with other teams are set up as copy members, and are used by both project teams. When a database or file layout change is made, usually all that's required is a recompilation of our system.

In addition to file layouts, we also use copy members for small (50 lines or less) frequently used intrinsic routines. *Program 1* shows an example of such an intrinsic, DBFIND. All of our database access is through canned, generic paragraphs in copy members. We have members for DBOPEN, DBFIND, DBGET, DBPUT, DBUPDATE and DBCLOSE. This alone has been a tremendous benefit for our team, since our I/O is standardized and programmers simply set up the parameters for the call and then perform these standard paragraphs.

We have similar copied paragraphs for all our V/PLUS calls. Again, terminal I/O is standardized; programmers simply set up parameters and perform the V/PLUS paragraphs. A side benefit of these standard paragraphs is that new hires who aren't familiar with IMAGE or V/PLUS can join our team easily, since they're isolated from the intrinsics themselves. Most important, however, is that we know these modules are bug-free and we don't need to reinvent them!

Our use of $INCLUDE files is usually for the portions of our system written in PASCAL, which cannot access our COPYLIB. (*Program 2* shows an example of a $INCLUDE file used by a PASCAL program.) We also use $INCLUDE for data areas required by modules supplied by other projects. For example, another project team provides a series of tables of valid locations, suppliers, code values, etc., and they supply an access routine that is stored as a module in our account SL. In order to call this routine, we must pass it certain parameters in linkage. We use a $INCLUDE file provided by the other project to set up the linkage area to this routine. We have similar files for each such module we use.

For modules shared across applications, we call routines in an account SL. Since several of our marketing applications reside in the same account for production, this makes sharing code very easy! For example, access modules for other applications' databases are located in the account SL. This eliminates each

application that reads a database code having its own access routine. It also ensures that database changes are transparent to those applications.

Other routines located in the SL include date routines, security verification routines and other utility routines. Since these routines are accessed at runtime, changing them is very easy and doesn't require updates to the calling applications.

If you are fortunate enough to be starting from scratch, seriously consider how to share code . . .

One word of caution about SL use: If a bad version of a routine is placed in an account SL, the results can be disastrous. For example, if many applications in one account all use one date routine residing in an account SL, and the date routine has a bug that causes aborts when validating February 29, you can imagine the havoc that will ensue on the morning of February 29 when all the applications blow up!

For this reason, our department has enacted strict testing and control procedures on routines placed in the SL. Changes to these routines must be tested in a group SL specific to one application before they can be moved into the account SL. This ensures that only bug-free code is accessed by all applications. Of course, all shared code should be thoroughly tested. However, SLs can be changed without the calling programs even knowing this has happened. For this reason they should be very carefully tested before being implemented.

The result of using these techniques is that a typical program in our application is 55 percent copied code from $INCLUDES and COPYLIB members. That's 55 percent tried-and-true bug-free code! More important, the programming time to make one change to one file and run a recompile and PREP is much less than changing all 75 of our individual source files, not to mention being much less error prone. It tends to save our programmers' sanity as well. In addition to code copied into source files, approximately 15 percent of our COBOL subprograms are called by more than one function within the system. This means 15 percent more proven, bug-free code.

HOPEFULLY THIS DESCRIPTION of one project's solution to shared code will help you implement similar techniques in your own project. If you are fortunate enough to be starting from scratch, seriously consider how to share code before you write a single line of source. (You may find the chart in *Figure 1* helpful as you set up your project.) This up-front effort will pay for itself over and over again.

If you are in the more typical situation of maintaining existing code, consider pulling out common code into $INCLUDE files, COPYLIB files or subprograms. Although it may be difficult initially, this work will pay for itself very quickly in reduced maintenance and fewer defects.

MIS organizations are under pressure to reduce costs and improve productivity. Using one or more of the above code-sharing techniques can improve your development environment on the HP 3000 and bring it closer to the ideal of a single point of change. It can reduce costs and reduce tedium for MIS staffers. Most important, reusing code will improve the quality of your MIS applications. That can only mean happier users. — *Lisa Burns Hartman manages an internal business software programming team at Hewlett-Packard corporate headquarters in Palo Alto, CA.*

Would you like to continue to see articles on this topic?
Enter On Reader Card
yes 812 no 813

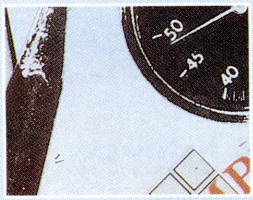

BUSINESS

Don Person

Tips On Purchasing Previously Owned, Reasonably Priced HP Computers

Confessions Of An HP Horse-Trader

Most computer manufacturers don't like to think about the back end of production processes or last year's equipment. HP is no exception. If you can get along without the newest and best, there's a good selection of nearly new or previously owned HP equipment for sale cheap all over the country.

Considering the reputation for quality that goes into HP computers, there may be more value in buying slightly used HP equipment than just about any other brand. Extensive use of gold connectors and top grade components makes a difference in overall reliability. The only exception to this might be in the smaller disk units that, for the most part, use the same modular components as every other outfit in the business.

For the past year and a half, I've been investigating and dabbling about in the used HP computer business, partly for profit but mostly for fun. If you don't mind a scratch, dent or blemish, you can get almost any older computer or peripheral somewhere for probably less than half the new price. With a little digging, you can get the hardware for much less. Here's what to look for, what to expect and how to bargain.

If turning surplus HP computer equipment into fun and money appeals to you, consider this: A good way to lose your shirt in any business is not to know the products you deal in. If you decide to play the HP used computer game, pick a model you know and stick to it. You might have to provide support for used gear, so this makes sense, doesn't it?

As in any other endeavor, some old advice still holds a lot of truth: Buy low and sell high. HP 120 series CP/M units, HP 85As, original 150As and single-sided 8″ disk drives are deader than Triceratops. If you have one, chances are nobody wants it for more than five cents on the dollar. That's the experience I've had and it's echoed by others who work this field full time.

For a wee bit of legwork, you can be rewarded with some great HP equipment . . .

Suppose you want to find a bargain, but don't know where to look or how to haggle for a good price. In most foreign countries and places that sell used equipment, the seller expects that you'll dicker over the price and terms of sale. Usually the used equipment dealer is surprised if you don't try to wheel and deal a little.

The sole exception to this might be your local HP consignment salesperson. Even he can see reason near the end of a sales quarter, especially if the goods you want are still sitting on his shelf.

In dealing with HP consignment staff, remember that they have a price floor too. The usual offer is 50 percent off, but 75 percent isn't out of the question. The formula is complex, but prices that are discounted 60 to 70 percent are often possible.

Just don't expect this kind of deal on units that are still the subject of active corporate promotion. For example, you want a LaserJet, but you don't want to pay more than $2,000. Now that the new LaserJet II models are out, some of the heat is taken off the original unit and the PLUS. Consignment people are likely to offer you a bargain price on one of these before they'll consider a cut rate offer on the new stuff.

There are always developers with surplus gear, although they're less available. HP sells systems at a 45 percent discount to software and other applications originators. Expect to pay 20-30 percent of the list price for their used equipment.

Don't rule out HP employees and OEM purchasers of HP equipment either. HP employees get discounts ranging from 40 to 65 percent on personal purchases. If you know a computer toy fanatic, you could find a real deal: deep discounts and light use. It won't get much better than that.

OEM purchasers of HP systems are another super resource. Usually they buy in quantity and sell surplus peripherals and computers in lots. One outfit I discovered had Vectra model 25s at 85 percent off. You just have to look around.

A surprising source of HP units is the small town newspaper or regional computer/electronic surplus firm. If you're lucky enough to be near an HP company town, you can find HP employees and their associates who have HP goods to sell. In the San Diego area, the *Byte Buyer* is a perfect example. A monthly publication, nearly every issue has classified ads placed by owners willing to sell HP personal computers and peripherals up to 90 percent off. There

are other such publications too. I used to scan several small computer publications to see who didn't sell their stuff and then make a late low offer. It works about half the time.

Remember educational institutions. HP's past practices have left computer equipment from micros to minis in schools and colleges. HP had a fairly extensive grant program for high schools, and more and more of the micro systems placed that way are ending up in the surplus market. Colleges and universities are in the same situation.

Then we have business bankruptcies. A regional merchandiser in my town had a nice 3000 system with 12 terminals, disks and tape that sold for about $8000 at an auction. You may have the advantage in making a lucky competitive bid, since in spite of HP's name recognition ad campaigns of recent years, the public at large has little brand name awareness. You can prosper from your knowledge.

Other industries out there are severely hurting. About one year ago, surplus computers began to appear as the oil and gas exploration industry in the US took a dive. Use your own imagination to target businesses that use HP systems and might be short of cash; then watch for auctions.

Healthy businesses also sell off excess HP gear. Surveyors, the research and development arms of major corporations and engineering firms are the best places to look. Larger companies and labs often have specific employees who handle this disposition. Call them.

Also consider institutions. HP's reputation as a medical equipment manufacturer has helped it make some hefty sales to a variety of hospitals and other health care facilities. When they sell, they can be an exceptionally low-priced resource. If you have the contacts, be sure to examine this possibility as well.

There are other advertised opportunities that may tempt you, but be careful. Less appealing to me are surplus units dumped by electronic instrument leasing companies. The computers that they sell may represent a much greater risk than other sources I've mentioned, and I personally consider them the most dangerous surplus to be involved in from the resale point of view. In nine cases out of 10, you'll find that the unit is being dumped because it's:

- *way out of date (unrentable)*
- *too beat up physically to be rented again*
- *unreliable or intermittent*
- *or all of the above.*

Since the price charged is usually higher than 50 percent of list and not subject to negotiation, I consider this particular source of used goods the most treacherous and least desirable. Will it be working? Probably. Cosmetically good enough to sell again? Unlikely. You'll do better with run-of-the-mill electronic surplus dealers.

I was shocked, for instance, to walk into the notorious back room at Eli Hefron's in Cambridge, MA, and behold tons of HP 3000 components from system units to peripherals. The prices? Since Hefron's is mainly a DEC surplus outfit, it was strictly, "Make me an offer." You can find some truly outrageous deals this way! Seek out the company not regularly reselling HP and cut a deal.

Now that you know where to look, what should you pay for a used HP computer? The real values are hard to fix, but here's the rule I've used over the past few years. With the exception of the real dead meat I mentioned earlier, 10 to 15 percent is a good price point to begin negotiation. If you pay more than 50 percent, chances are that you failed to find a direct consignment deal from HP and are paying too much. Remember, for a wee bit of legwork, you can be rewarded with some great HP equipment at a real savings. Happy hunting!
—Dave Person is an independent consultant based in Albany, NY.

Would you like to continue to see articles on this topic?
Enter On Reader Card
yes 814 no 815

Computer Fraud Upon Our Children

OPINION

Cliff Lazar

Many young children go to schools where LOGO is taught as the first computer course. Parents who are computer professionals or lay non-computer people are told that their children are receiving valuable computer instruction.

Actually, these children are being given meaningless and counter-productive LOGO exercises by a teacher who is otherwise not competent in computer technology.

MY INSPIRATION FOR WRITING this article came when I was in a book store and overheard an elementary school teacher tell her friends that she was teaching LOGO. As she understood it, LOGO was like giving instructions to a turtle to draw lines on the screen that would combine into simple figures to make pictures.

Yes, the instructions would combine to make simple figures, if the instructions were perfect. Otherwise, the figures would look funny and everyone could laugh at your child's efforts. Forget that the figure assignments were little stick houses and stick people. But don't forget that the keyboard time usually is inadequate and the feedback isn't instantaneous.

The assignments, like all programming, require debugging, but your children get inadequate or no keyboard time. Lack of perfection doesn't result in nurturing, but in downgrading. The linear thinking and the perfection required to make perfect LOGO figures is an unusual skill that's in very low demand in our modern spell-correcting society.

Programming computers is a process of creative and reflective problem-solving. Sometimes errors take a while to debug. Programmers, like those who created the MacIntosh interface, Visicalc or Pagemaker, weren't forced to be perfect the first time or to have a high rate of keystrokes.

LOGO has many opportunities to have failure experiences, but only one way to create what the assignment specifies. Good educational practice involves success experiences that reinforce each other.

LOGO is the ideal opportunity for ill-prepared teachers to act like pedants. According to Webster's: "*ped-ant* 1. one who emphasizes trivial points of learning 2. a narrow-minded teacher who insists on exact adherence to the rules . . . ." As such, they exercise clearly improper power over clearly imperfect children. LOGO lends itself to such pedantry because its exercises are trivial and can't create anything of real use or beauty.

When I took print shop, the tech-

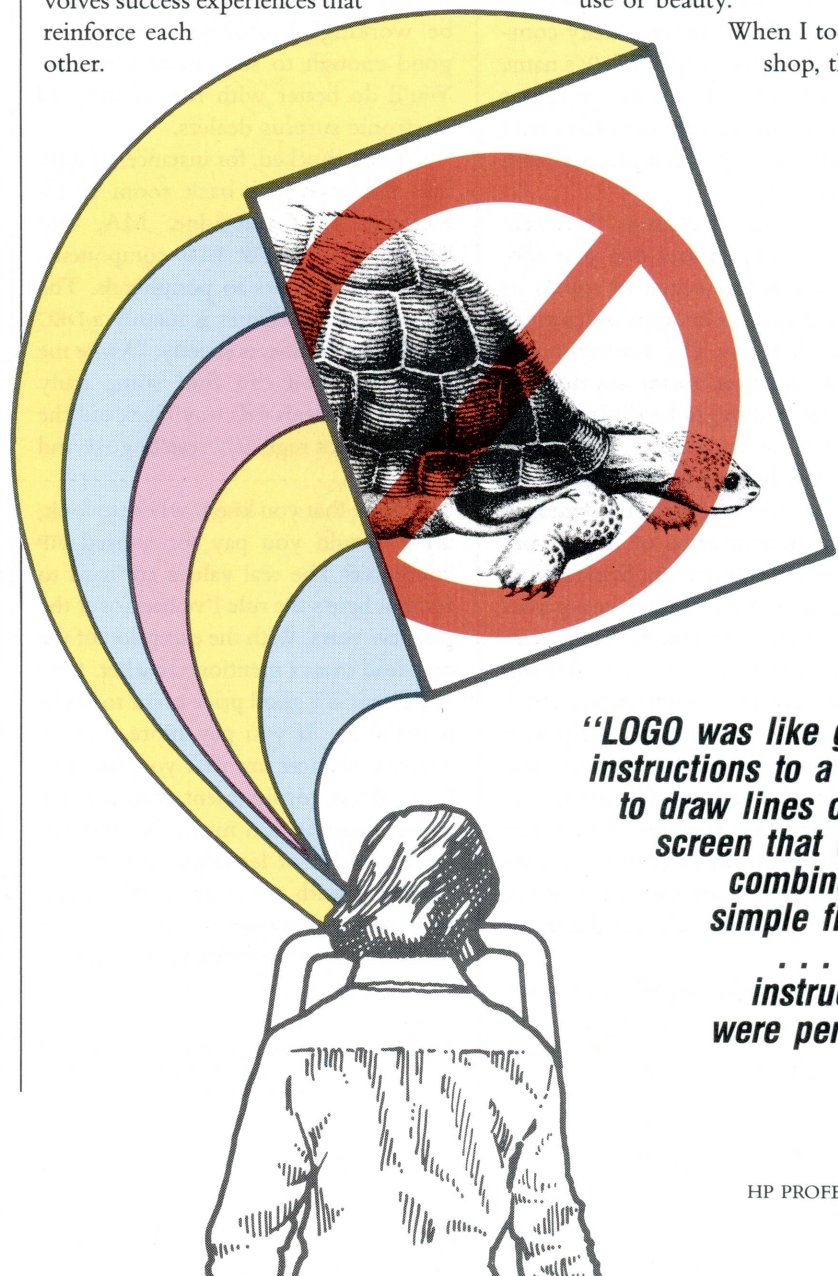

"LOGO was like giving instructions to a turtle to draw lines on the screen that would combine into simple figures . . . if the instructions were perfect."

nology was already obsolete. We were setting cold type when the world was using linotypers and lithography. LOGO is similar. The interface is a keyboard with many keystrokes required for very limited results.

Today's technology for graphic input is a mouse. The hand moves and the

The linear thinking and the perfection required to make perfect LOGO figures is an unusual skill that's in very low demand . . .

computer does the work. LOGO requires that the student act like a computer. If you want to draw pictures, you can use computer-aided drafting tools that generate very useful drawings and cost less than LOGO. If you want to do math, process information or control machines, you can use BASIC, which comes free on all Apples.

I taught a paraplegic how to use a computer-aided drafting system, and in 45 minutes he was generating drawings that equalled the quality of two-handed professionals.

If I had had to teach him LOGO, he would have quit out of frustration in the first 20 minutes. The drafting system internalized the LOGO-like plotter commands. He didn't have to care about each command to push around the plotter pen, which is what LOGO is about. All he had to care about was what his drawing would look like. The feedback was instantaneous. He'd point at both ends of a line and it would appear. He'd

specify a circle with a center and an edge and it would appear — in two keystrokes. If it were in the wrong place, a single undo command or a move would fix it.

There is a better alternative for technology training. First, understand that LOGO is a surrogate stepchild for technology in general. "Don't worry, parents, we're teaching your children LOGO." I feel we should teach our children about technology: Where it was, where it is and where it's going. Computers are only one part of modern technology. We also should teach about communications and automation.

We should understand that more people will use computers than will program, build or fully understand them. We all know that people who feel comfortable with computers, because they're useful and fun, will have an advantage over those who were turned off by pedantic LOGO teachers. I guess that LOGO will be to computers what the metric system is to science and math. Well over half the American public who were trained in the metric system don't understand it and can't use it. The same will be true of turned-off LOGO victims.

MOST PEOPLE DRIVE because it's simple and very productive. For teenagers, it provides an opportunity to get privacy, and what possibilities can stem from that. Computers can be simple and immediately rewarding, too. And that's what we should teach in school.

Have them do their homework on the machines. Let them word process their essays and book reports. Let them graph their math, sociology and science. Have them use a spreadsheet like Lotus 1-2-3 or Visicalc to collect data and do totals. Show them how to build a check register and compute baseball averages and budgets. These are the tools they'll need in later life — not how to give micro commands to a plotter pen. They also should be taught computer survival skills, such as plugging together the computer, disk drives and printer. When to format and not format disks, and what it means, should be taught.

Computer classes can demystify what business people and college students do, and prepare children mentally for exciting opportunities and adventures in business, education and government . . .

Teach them that computers can communicate. Teach them about modems and baud rates, and about all the information that exists in Compuserve and in databases. School children in San Diego transmit pen pal letters to students in Alaska. Your child could send EasyLink messages or telexes like the world of business does.

Computer classes can demystify what business people and college students do, and prepare children mentally for exciting opportunities and adventures in business, education and government, instead of having them pretend they're little computers themselves.

We should tell children what computer technology is all about and that attitudes change and people learn as technology changes. We should teach them that at one time everyone built VonNeuman model single-CPU machines, but that parallel computers are taking over.

OUR CHILDREN SHOULD LEARN that computers can make decisions and that the Star Wars Battle Management computer systems can make mistakes.

They should learn that computers are going to be used more and more for real-time control. They should learn that assembly line jobs are going to be displaced and that honest manual labor in a factory may become white hat labor in McDonalds.

Our children should learn that computers can be programmed to act like human — albeit narrow-minded human — experts. Acting like knowledge engineers would be a good exercise. Collecting the decision process that their teacher uses to evaluate their grade performance or what the coach uses to pick a football play, or how the school board picks a school site would be a good exercise.

Some, but not all, of our kids should get to do advanced programming, like some should get to do advanced chemistry or physics. BASIC is the best start. It should be taught so that the programs are structured and not spaghetti-coded. It's not clear that the next language should be PASCAL. You can write unstructured PASCAL and you can write structured BASIC.

Some schools actually say LOGO should be taught before PASCAL. That's like saying you should ride a horse before you can drive a car. Having programmed in both, I promise you LOGO is not junior PASCAL. LOGO may have procedures, as PASCAL does, but BASIC has GOSUBS which are the same. LOGO doesn't have data structures or records or meaningful results; BASIC has all of them.

Since the market is still big for COBOL and there now are low cost COBOL compilers for PCs, maybe a subset of COBOL should be taught to those who are interested in business problems. PASCAL, C and FORTRAN can be taught to those who are into science and math. But *always* the teachers should understand that *no single way is the best for every child.*

TEACHERS WHO DON'T dance shouldn't teach dancing. Teachers who don't program shouldn't teach programming. Programming is a way of thinking that can't be taught by someone who doesn't think that way.

When I went to school, no one thought the English teacher qualified to teach metal shop after a two-week course. You needed experience with metal and ovens. No one thought the gym coach could teach home economics after a two-week course. Why do educators think computer programming is so much simpler?

Education is passing knowledge and skills from those who have them to the next generation. It is not going to a two-week in-service educational enrichment class and then inflicting a lack of knowledge, background, enthusiasm and skill on a powerless child.

If you, as computer professionals, feel strongly about this, you can do something. Send this article to your child's principal along with your personal comments. Since your principal may be a graduate pedant who has taken enough educational enrichment courses along with developing bureaucratic skills to stop teaching and become an administrator, he'll likely reject or ignore your comments.

This will allow you to take the next step — send the article to your local school board member along with your comments. Hopefully it's an election year and he'll be more sensitive.

A next step would be to send your comments to the local newspaper education editor. If he receives as many as five, he'll react and do a story — which will open a can of worms and begin a debate that could keep your children from being turned off by your profession.

It's a worthwhile endeavor. You wouldn't want your children thinking that you spend your day instructing turtles. —*Cliff Lazar is president of Systems Express, Sherman Oaks, CA.*

Would you like to continue to see articles on this topic?
Enter On Reader Card
yes 810 no 811

Continued from page 25.

HP, DEC and IBM networks; intended use by medium- to large-sized companies that have installed HP 3000s, DEC VAXs and IBM mainframes. These companies are then able to integrate their computer resources, making the data and application software on each computer accessible by users of the other vendors' computers.

GATEWAY/1000's support for tri-directional virtual terminal sessions allows any HP, DEC or IBM terminal to access any computer. Interactive access from HP and IBM terminal networks to a DEC VAX computer is provided by DECnet and VT220 terminal emulation. Interactive access from HP and IBM terminal networks to a DEC VAX computer is provided by HP 2622 terminal emulation. Interactive access from HP and DEC terminal networks to an IBM mainframe is provided by IBM SNA 3270 terminal emulation.

GATEWAY/1000 hardware and software prices begin at $18,000 and $16,000 respectively.

Contact Forest Computer Incorporated, 1749 Hamilton Road, Okemos, MI 48864; (517) 349-4700.

Enter 906 on reader card

MONTEST-A4: New PC Video Generator

The MONTEST-A4, a new PC video generator, provides low-cost equipment for graphics personnel, CAD/CAM and other software designers. Battery operated, it directly drives analog style monitors including the PGD professional graphics monitor. The unit drives monitors directly providing a range of switchable horizontal scan rates from 15.75 kHz to 31.5 kHz.

The MONTEST-A4 displays four separate patterns in normal and reverse video including color bars, cross-hatch, window and full raster with individual color and intensity switches. It also allows a separate sync output or composite sync on green. The user can test and set color purity, convergence and alignment.

The MONTEST-A4 is available from stock at $595 including AC adapter and one-year warranty.

Contact Network Technologies Inc., 19145 Elizabeth Street, Aurora, OH 44202; (800) RGB-TECH or (216) 543-1646.

Enter 919 on reader card

Unison Launches Spoolmate

Spoolmate is a new data center management tool that provides extensive report production and central console control of all spooling activity, control of multiple-CPU spool files and the availability of multiple queuing levels.

Operators have the ability to control printing at the device level; i.e., reserving printers based on spool-file size or stationery type. Automating printing of spool files according to stationery type means less operator time spent changing paper. Assigning spool files according to individual printer capabilities means increased throughput due to increased efficiency. Device-level control is a new concept in spooling and one that will allow users time-saving capabilities unknown in today's market.

Other features of Spoolmate include: complete audit trails of spooler activity, report distribution tracking, archiving spool files to disk or spooler, and diagnostic capabilities.

Spoolmate is value-priced according to CPU size, and is available for the Spectrum. Contact Unison Software, 415 Clyde Avenue, Mountain View, CA 94043; (415) 968-7511; and Europe via Telex 928536.

Enter 907 on reader card

Koch Software Introduces DwLaser

DwLaser is a QuietWriter II emulator, which supports all Displaywrite features including margin control, headers/footers, pagination, laser- supported symbols, font support, and right justification with proportional fonts with any HP laser printer.

DwLaser is a system package that allows users to choose command options that specify fonts, printer, orientation, pitch and lpi adjustments. All option choices are saved to a batch file which is automated each time Displaywrite is used.

The Koch Symbol Fonts included with DwLaser contain all symbols that are accessible in Displaywrite — available in 10-, 12- and 15-pitch landscape and portrait. Users also may combine the Koch Symbol Fonts with other HP cartridge and soft fonts. DwLaser is retail-priced at $195.

Contact Koch Software Industries, 11 West College Drive, Bldg.-G, Arlington Heights, IL 60004; (312) 398-5440.

Enter 908 on reader card

Versatec Offers HPGL Plotting Utility

Versatec has released a new HPGL Plotting Software Utility that converts Hewlett-Packard Graphics Language (HPGL) commands into Versatec Versaplot software commands for output to Versatec plotters. This utility supports all Versatec electrostatic and thermal transfer color plotters.

The Versatec HPGL Plotting Software Utility is available and integrated for the IBM AT personal computer and must be used with a Versatec Model 110 interface board. The utility performs the ordering, rasterization and output to the driver and then the plot-

MONTEST-A4—Network Technologies' newest PC video generator.

ter. It is menu driven, making it easy to select options such as changing pen widths and selecting colors.

The HPGL utility will accept the most commonly used commands in HPGL application software. These commands will be converted by the utility into Versaplot software commands for output to Versatec plotters.

Versatec HPGL Plotting Software Utility is available for the IBM AT PC priced at $100 on floppy disk. Contact Versatec, 2710 Walsh Avenue, Santa Clara, CA 95051; (800) 538-6477, in CA: (800) 341-6060.

Enter 911 on reader card

HPS Introduces Termtype

Using existing Microsoft Windows software, HPS Software Developments Limited has developed the new Termtype terminal emulator, which imitates the normal functions of an HP 2392A terminal and can display applications that are running under Windows at any one time on its monitor.

Termtype supports all the normal functions associated with Windows such as "cut-and-paste" and "data-sharing," thus providing a versatile user interface for MS-DOS-based machines.

Contact HPS Software Developments Ltd., 196A Whittington Road, London N22 4PD UK; Telephone 01- 881 6644.

Enter 909 on reader card

Speededit Passes Preliminary Testing

Bradford Business Systems has completed preliminary testing of Speededit on the HP 3000/930 Spectrum System and reports no problems.

Speededit is a program development system capable of true full-screen editing, inline compiles, version tracking, multiple simultaneous edits on multiple files, multiple open edit sessions with suspend/activate feature, and over 100 editing functions. The system can be user-customized to better suit programming style and is also useful for: general text editing, electronic mail, tickler files and built-in spell check/corrector.

Contact Bradford Business Systems Inc., 25301 Cabot Road, Laguna Niguel, CA 92677; (714) 859-4428.

Enter 910 on reader card

AP500 Available For The HP 9000 Series 800

The AP500 Array Processor, a hardware and software package, has been developed to interface the AP515 Full Floating Point Array Processor to the new HP 9000 Series 800 Precision Architecture minicomputer systems. The first of these new systems is the HP 9000, Model 840, running HP-UX, a super set of AT&T System V interface definition (issue 2).

The processor interfaces to the Series 800 via two HP CI/O interface boards, occupying two consecutive CI/O backplane slots. It supports Programmed I/O (PIO) and Direct Memory Access (DMA) data transfer between the HP computer and the array processor. The measured average DMA transfer speed is 2.2 MBs.

The software package includes the device driver, a library of more than 350 array processing functions callable from

FORTRAN and C, and utilities for extensive programming development and diagnostics. Single units start from $32,000.
Contact Analogic Corporation, 8 Centennial Drive, Centennial Industrial Park, Peabody, MA 01961; (617) 246-0300.

Enter 913 on reader card

Office Extend For HP 150 Series PCs

Like Office Extend for IBM-compatible PCs, the HP 150 implementation provides a transparent print path from the Touchscreen to HP 3000 printers, allowing application programs on the 150 to treat spooled host printers as if they were directly connected to the Touchscreen.

Office Extend will work with existing 150 to HP 3000 connection schemes and no hardware or operating system upgrades are required. It is compatible with almost all software running under DOS 2.X or higher, and consumes less than 8K memory on the Touchscreen while it includes a 2K print buffer. The price starts at $1,245.
Contact Fransen/King Ltd., 16400 Southcenter Parkway, Seattle, WA 98188; (206) 575-1570.

Enter 914 on reader card

IBM Proprinter Emulation Added To LaserPro Line

The upgrade of the IBM Proprinter Emulation as a standard feature on all LaserPro models has been announced by OASYS, with no additional cost to the end user.

The Proprinter emulation supplements standard emulations currently available with LaserPro printers, including: HP LaserJet/LaserJet Plus, Diablo 630, Epson FX-80, NEC Spinwriter, Qume Sprint 11 and ANSI 3.64.
Contact Office Automation Systems Inc. (OASYS), 8352 Clairemont Mesa Boulevard, San Diego, CA 92111; (619) 576-9500.

Enter 915 on reader card

Datacaptor IV Enhanced

A plug-in circuit board providing enhanced capability, including alphanumerics for bar code scanning devices useable with the Datacaptor IV Point-of-Sale (POS) terminal, is available for the following units: Symbol Technology laser scanners (Models LS-7000 and LS-8000), HP pencil-type wands and Metrologic table top scanners.

The enhancement board is priced at $325 for the laser scanner configuration, $575 for the handheld configuration (including scanning wand) and $1695 for the table top configuration.
Contact Datacap Systems Inc., 212A Progress Drive, Montgomeryville, PA 18936; (215) 699-7051.

Enter 916 on reader card

Outcome Provides Sales Information

Designed to increase productivity of a sales force, Outcome offers customer data management, past contact records, sales objectives, decision-makers involved per sale, ideal customer match, forecasts, sales funnels, electronic mail and many other features.

Acknowledging that the typical salesperson avoids any sales tools that require extensive training or rely on complex documentation, an intuitive interaction is created between the software system and the salesperson. Knowledge of computer commands or database operations are unneeded.

The software may be used on IBM PC compatible laptop and desktop computers, and HP 3000 and DEC VAX minicomputers. Outcome is priced at $500 per salesperson.
Contact Profit Management Systems Inc., 9800 Fourth Street North, St. Petersburg, FL 33702; (813) 578-0190.

Enter 912 on reader card

CDS/SecReview: Security Management Tool

Chestnut Data Systems has announced version D.01 of CDS/SecReview with several new features including MPE password aging.

CDS/SecReview, a system security management tool, reports the degree to which MPE system security features are utilized.
The package is priced at $749 and includes documentation and one-year free maintenance. Contact Chestnut Data Systems, Park Towne Place, S-05, 2200 Benjamin Franklin Parkway, Philadelphia, PA 19103; (215) 557-6607.

Enter 918 on reader card

OCS Announces OCS/EXPRESS

Operations Control Systems (OCS) announces OCS/Express, a new product in the OCS line of data center automation software that builds on the batch scheduling capabilities of OCS/Scheduler and OCS/Dispatcher. OCS/Express' new utilities that automate installation activities and schedule entire sets of jobs with a single command.

BBS-PRISM

FULL SCREEN EDITING FOR THE TERMINALLY SLOW

See you at Interex Booths #336, #338.

To run a healthy business, you need a computer system that won't keep your people waiting.

BBS-PRISM is a powerful full screen editor that lets you off-load just about all your editing functions to **PCs,** leaving the host system free for everything else. Your host won't get bogged down, and your people won't get frustrated.

We call this feature "Remote Program Development." You'll call it a lifesaver.

BBS-PRISM is the only rational approach to remote development, with its powerful features and the fact that only changes in the source file are transmitted back to the host, not the entire file. You get faster, smoother editing, with greater capabilities, in a package that is totally compatible with SPEEDEDIT, our 3000-based full screen editor and even EDIT/3000.

BBS-PRISM is loaded with useful features like built-in word processing.

Get all the details. Call Bradford Business Systems today. We've got an entire line of software designed to bring your business up to speed.

BBS-PRISM. A genuine contribution to office productivity.

BRADFORD BUSINESS SYSTEMS, INC.

25301 Cabot Road, Suite 201
Laguna Hills, CA 92653
(714) 859-4428

ENTER 194 ON READER CARD

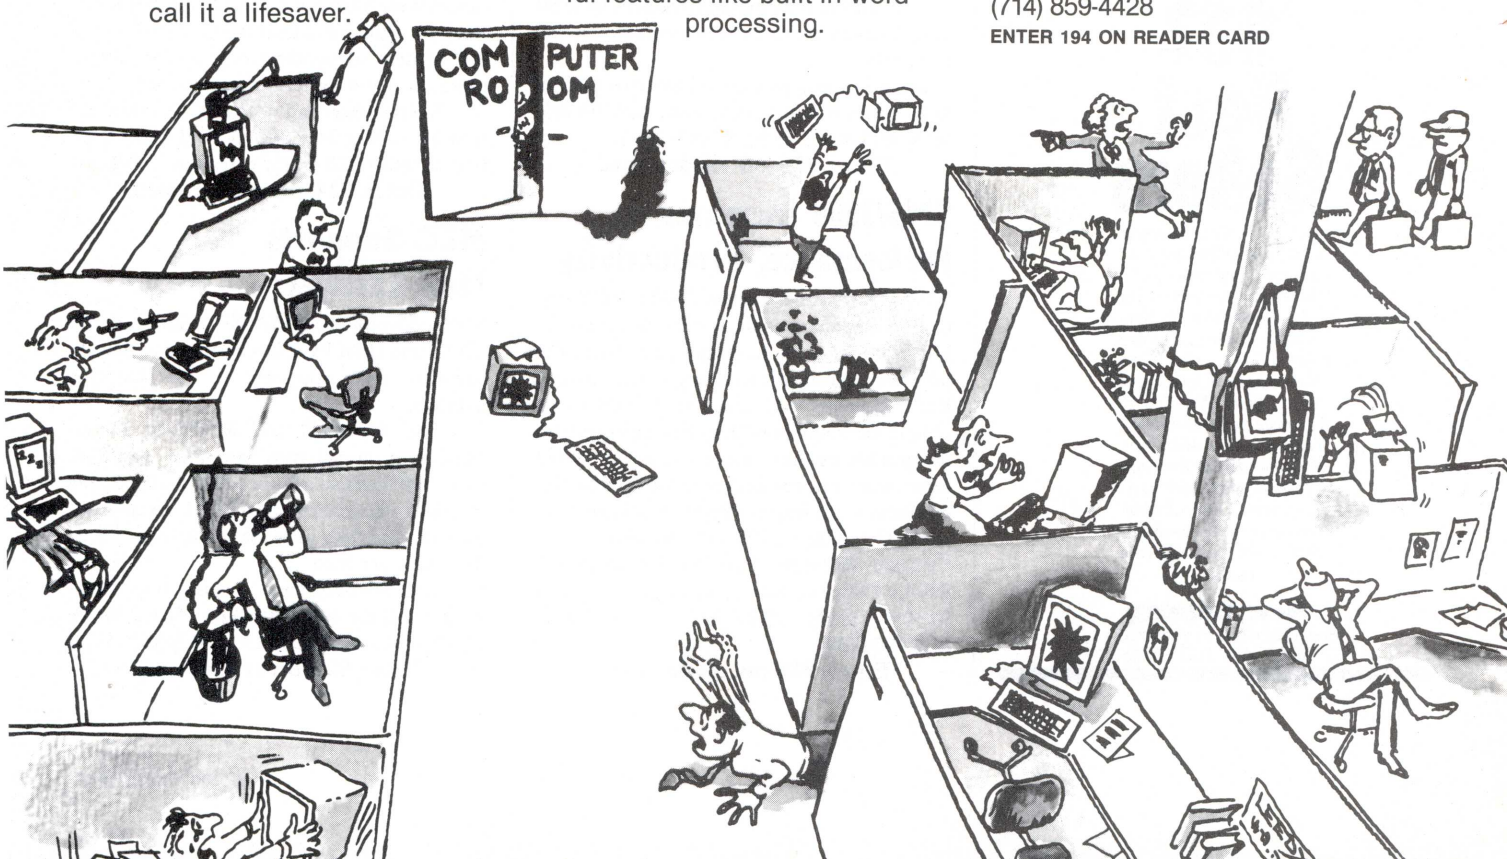

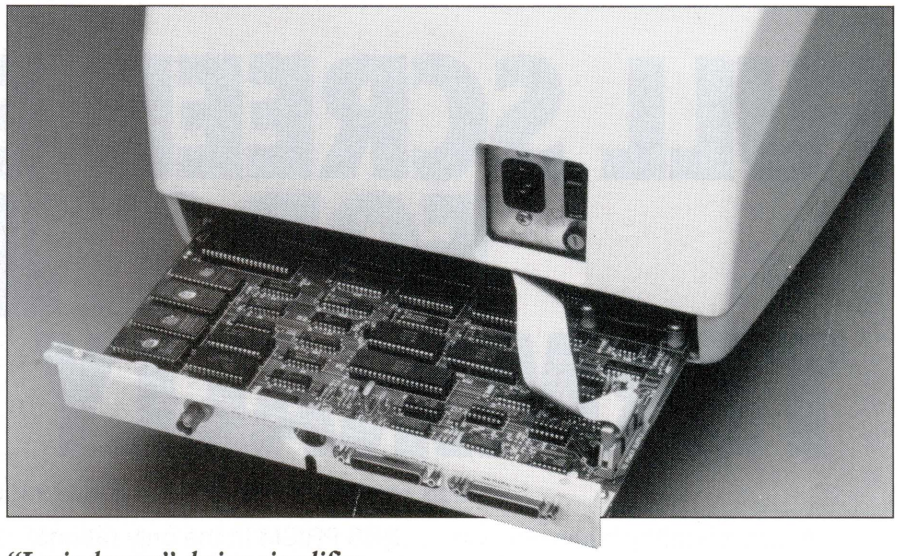

"Logic drawer" design simplifies servicing of Teleray terminals.

OCS, an active participant in the Fast Start vendor program, also announced the availability of OCS/Express on the Series 903 and 950 systems.

OCS/Express allows users to define their own custom calendars to determine when jobs will run. Additionally, an unlimited number of prejobs, preschedules and preevents can be defined in one step. The system also provides selective displays of production status as well as immediate job processing.

In all, OCS/Express introduces over 50 new features previously unavailable on the HP 3000.

Contact Operations Control Systems (OCS), 560 Antonio Road, Palo Alto, CA 94306; (415) 493-4122 or FAX 493-3393.

Enter 920 on reader card

SMS/3000 Increases Performance, Productivity

The Sales Manager System/3000 (SMS/3000) is a sales management system designed for HP 3000 users. It manages leads from the point of identification through sales closure. Each company can adapt SMS/3000 to its unique marketplace with its five independent, integratible modules which can be purchased separately or in combination: Sales Manager, Business Card Index, Project Tracking, Personal Calendar and Electronic Mail.

Contact Business Systems International, 20942 Osborne Street, Canoga Park, CA 91304; (818) 998-7227 or TELEX: 51011004239.

Enter 921 on reader card

Teleray Terminals Receive Maintenance Services

Teleray has established a maintenance agreement with Honeywell-Bull, providing service to four Model 20 multiprotocol terminals.

According to the agreement, the 20 DDG (D211/DEC VT220 compatible), 20-DHP (HP 2392A/DEC VT220 compatible), the 20-7305 and 20-7813 (both Honeywell/DEC compatible) will receive onsite with four-hour or next-day response time, or walk-in/mail-in at Honeywell service centers depending upon the user's preference and system requirements.

Teleray users can obtain more information by calling Honeywell-Bull CSD toll-free at (800) 328-5111, extension 2769.

Enter 924 on reader card

MEC Supports New LaserJets

MEC, developer of MASS-11 Software for VAXs and IBM PCs, has announced that its software now supports all HP LaserJet printers.

The MASS-11 and MASS-11 Draw, MEC's mouse-driven graphics package, gives HP users an advanced level of text and graphics integration. MASS-11 Draw outputs to LaserJets in resolutions from 75 to 300 dots per inch.

Contact Microsystems Engineering Corporation, Suite 400, 2400 West Hassell Road, Hoffman Estates, IL 60195; (312) 882-0111.

Enter 922 on reader card

STARNET, RQM Form Alliance

Denniston & Denniston Inc., and Automated Technology Associates have formed an alliance for the development and marketing of integrated factory automation systems. The two companies have integrated their two principle products: STARNET and RQM on the HP 1000 family of real-time computers.

STARNET, of Denniston & Denniston Inc., is a supervisory control software system used for Computer Integrated Manufacturing. The software is a family of actual applications to fulfill most needs for "Level 2" monitoring and control, factory information systems.

These subsystems are user configurable for a variety of industrial control applications and production environments.

RQM (Real Time Quality Management), of Automated Technology Associates, is a real-time statistical process control system, capable of high-speed analysis of quality data, automated interpretation of quality control charts, and processing of statistical quality alarm conditions.

The integration of STARNET and RQM allows users to access each system directly from each other.

Contact Denniston & Denniston Inc., 3250 North Arlington Heights Road, Arlington Heights, IL 60004; (312) 398-8500.

Enter 923 on reader card

Safe C Family For HP UNIX Systems

Catalytix Corporation, a supplier of advanced software development tools for C programmers, has announced the availability of its entire Safe C product line for HP's UNIX-based computers.

Catalytix' Safe C family of software development tools consists of an advanced Interpreter for interactive execution of C programs and a Run-Time Analyzer for automatic detection of C programming errors. The Analyzer also includes modules that perform tracing, facilitate code optimization and analyze test coverage.

Catalytix also has released a special version of the C Trainer for HP's UNIX-based computers. The C Trainer is an interactive instructional package for teaching the C Language. It includes a textbook published by Prentice-Hall, *The C Trainer Interpreter*, and an online library of programs described in the textbook. The C Trainer will be sold worldwide in English, German and Japanese editions.

The prices for Catalytix software for

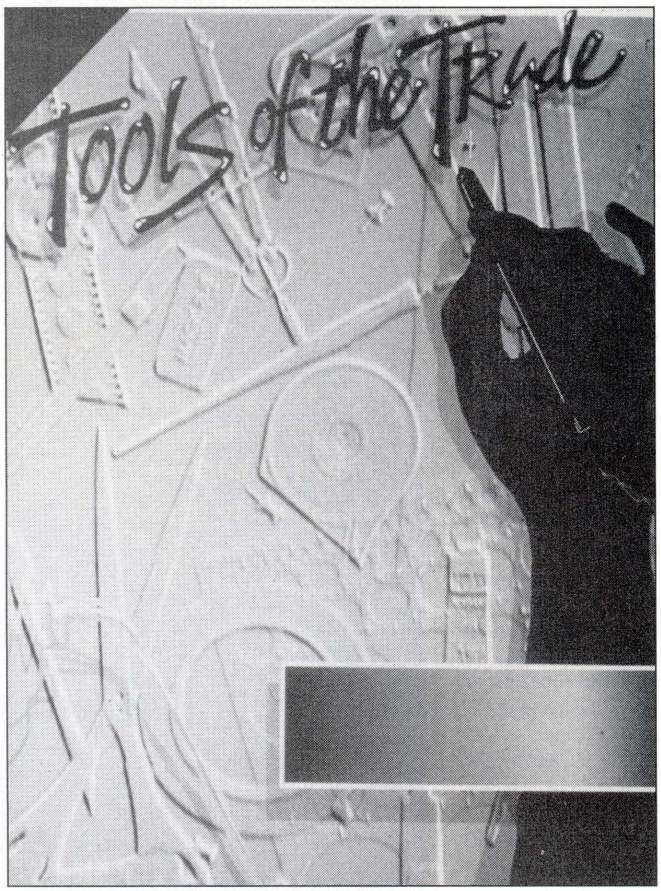

This image was created by artist Claire Barry using Lumena/16 and was output directly to a Scitex printer.

HP's 9000 Series 300 and 500 systems are: $2,000 for the Interpreter; $1,200 per Run-Time Analyzer module; and $220 for the C Trainer package. Prices for the Catalytix software on HP's new Precision Architecture Series 800 machines are $4,000 for the Interpreter, $2,400 per Run-Time Analyzer module and $420 for the C Trainer package. Contact Catalytix Corporation, 55 Wheeler Street, Cambridge, MA 02138; (617) 497-2160.

Enter 929 on reader card

Time Arts Enhances LUMENA Software

Layout grids and color filters have been added to LUMENA software.

The 10 major grid tools allow precise layout of images and text, and new options for design specifications. Perhaps the most innovative grid tool is the LUMENA version of GRAVITY, which allows the designer to select the distance from which the cursor is drawn to a given element on the grid. Users will also be able to create rulers, position them anywhere on the screen and set them to measure inches, centimeters and picas, further enhancing the precision of images created with LUMENA.

Of significance to the video market is the new Color Filters tool. Video professionals now can manipulate color balance, add tints to images and create highlighted windows.

Time Arts' product line continues to support input and output devices including digitizing tablets; input scanners; video cameras; thermal, ink-jet, dot matrix and laser printers; analog and digital film recorders; NTSC encoders; sync generators and animation controllers.

Contact Time Arts Inc., 3436 Mendocino Ave., Santa Rosa, CA 95401; (707) 576-7722; Telex 5106011663.

Enter 928 on reader card

New Device Drivers For Freelance Plus

Lotus Development Corporation has introduced new device drivers for Freelance Plus, the company's graphics product. They are Apple Laserwriter laser printer, Epson LQ-800/LQ-1000 dot matrix printers, HP

PaintJet color ink jet printer and CalComp ColorMaster plotter/printer.

Freelance Plus offers users graphics capabilities for creating business communications and adds another level of integration with Lotus 1-2-3 and Symphony allowing users to import data from worksheet files.

Lotus has also extended its Freelance Plus Upgrade program. Current Freelance program users can upgrade to Freelance Plus through December 31, 1987. The cost of upgrading is $100. Customers also may deduct the $100 upgrade price from the cost of any Freelance Maps package, a collection of Freelance Plus companion map sets, ordered at the same time. The maps cost $145 individually and $395 for the U.S. Complete set.

Contact Lotus Development Corporation, 55 Cambridge Parkway, Cambridge, MA 02142; (617) 577-8500 or (800) TRADEUP for upgrade kit information.

Enter 925 on reader card

Express Series II Emulates HP LaserJet Printer

OAS has announced the emulation of the HP LaserJet Plus, Express Series II. Express Series II offers HP font cartridge capability, 640 KB memory, and emulations including Epson America Inc.'s FX-80. Other features include: full-page 300x300 dot/inch resolution and resident RS-232 and Centronics Data interfaces.

The Express Series II is priced at $2,295. Contact Office Automation Systems Inc., (OASYS) 8352 Clairemont Mesa Boulevard, San Diego, CA 92111; (619) 576-9500.

Enter 926 on reader card

New LARC Laser Package Prints Without Cartridge

The LARC Laser Package prints bar codes without the need for a bar code font cartridge. The LARC facility is more flexible than a font cartridge. The user can determine bar widths/lengths to be used and also print bar codes in landscape orientation with text in portrait (or vice versa).

While the HP bar code font cartridge is temporarily off the market due to adjustments being made to suit the LaserJet II, the LARC capability can be used if bar codes are needed.

The Bar Code capability is the latest in a series of enhancements to the LARC Laser Package. The fee for software use and support is $2,250. Contact LARC Computing, 339 South San Antonio Road, Los Altos, CA 94022; (415) 941-9310.

Enter 927 on reader card

VESOFT Enhances SECURITY/3000

VESOFT, Inc., has enhanced its SECURITY/3000 logon access control package for the HP 3000 with the inclusion of a new Password History feature.

The new feature allows the system to maintain a history of previously used passwords for a specific time period (for example, one year). The history then automatically can be examined when it comes time to change passwords, in order to ensure that the new passwords were not used previously during the designated time period.

The new Password History feature is incorporated within one of SECURITY/3000's programs, which allows users to change their own passwords and thus relieves the system/account manager from the duty of changing each user's password.

Contact Eugene Volokh at VESOFT, 1135 South Beverly Drive, Los Angeles, CA 90035; (213) 282-0420.

Enter 931 on reader card

SystemsExpress Offers :DBMAINT Tool

:DBMAINT makes it possible to update rapidly and delete IMAGE records. SystemsExpress is adding the fifth generation IMAGE/3000 tool to its product line.

:DBMAINT is the result of user requests. It requires very little training and supports IMAGE/3000 data security. It isn't a tool for casual users, because it can corrupt a database quickly. Database administrators need to limit the release of the WRITE password to only those sophisticated programmers who can be trusted with the database. To increase the security, SystemsExpress has provided a password for :DBMAINT itself.

:DBMAINT is classified as a fifth generation tool because it has no language, only menus, and because it generates both third and fourth generation programs. Typical setup times will vary from five to 15 minutes. No classes or formal training is required.

:DBMAINT is priced at $2,000. Companies desiring a free test should contact SystemsExpress, (818) 784-6966.

Enter 933 on reader card

Dispatch, RSVP Added To Complements Line

A new product line will promote further integration of the PC as the complete workstation, announced the developer of Reflection Series terminal emulation software, Walker Richer & Quinn. The Reflection Complements are designed to take advantage of the PC as a workstation by addressing needs that are unique to the PC-to-mini connectivity.

The two new products in the Reflection Complements product line are Dispatch and RSVP (Reflection's Spooled Virtual Printer). Users of Reflection Complements never need to leave the Reflection environment in order to make effective use of host peripherals and office automation software.

Dispatch was created to maximize performance and simplify procedures in the HPDesk electronic mail program. Dispatch users can compose and read their HPDesk mail while offline from the HP 3000, and upload/download messages automatically at predetermined times. Dispatch also lets the user send spreadsheet, database or word processing files through HPDesk to other HPDesk users or host-based printers.

RSVP is designed to facilitate printing PC-generated printouts on a host-connected printer. Usually, the PC user must print to a file, upload the file, then print it on the HP 3000. The fact that some PC programs don't allow printing to a file introduces more complications.

RSVP makes printing from a PC to a host-connected printer a one-step process. Simply run RSVP from Reflection's command line, and all PC printouts automatically will be rerouted to the HP 3000-based printer. Source code is included so that RSVP can be modified to accommodate any device defined on the HP 3000.

More information is available by calling Walker Richer & Quinn at (206) 324-0350.

Enter 934 on reader card

Repeat-O-Type Offers New Cartridge System

Repeat-O-Type Manufacturing Corporation announces a high-yield replacement cartridge system for use in Apple, Canon, HP and other laser printers using Canon "CX" engines.

Approximate copy yield is increased from 3,000 copies to 15,000 copies per cartridge. Cost per copy is reduced from about 4 cents to 1⅔ cents. Suggested retail price is $259.95.

Contact Repeat-O-Type Manufacturing Corporation, 665 State Highway 23, Wayne, NJ 07470; (201) 696-3330 or out of state (800) 526-5200.

Enter 930 on reader card

Graphicus Products For HP 9000

Grafit, Drawit, and Stat80, software products from Graphicus, are now available on the HP 9000 Series 800, HP's family of Precision Architecture computers. All three Graphicus products have a large installed base in the HP 1000 user community.

Two of the products, Grafit and Stat80, also have been implemented on HP's desktops, on the HP 9000 Series 300 and 500 under UNIX. Availability of Drawit on the desktop will be announced soon.

Tailored for the scientific and technical user, Grafit, Stat80 and Drawit together provide an integrated data analysis and data presentation capability. Each also is a standalone system. All three are HP + referenced products.

Contact Graphicus, 160 Saratoga Avenue, Suite 32, Santa Clara, CA 95051; (408) 246-9530.

Enter 932 on reader card ■

Correction

In Vol. 1, No. 1 (May 1987), incorrect pricing information was given for the HP ScanJet and the bidirectional parallel interface card and *Scanning Gallery* software (p. 24, "New Products"). The correct prices are:

ScanJet alone — $1495.

Interface Card and *Scanning Gallery* software — $495.

If You Are Considering HPSQL...

Consider This:

Function / Feature Comparison	HPSQL	RELATE / 3000
Fully Relational	Yes	Yes
Accesses Existing Image Files	NO	Yes
Integrates With A Report Writer	NO	Yes
Integrates With A Graphics System	NO	Yes
Integrates With A 4GL	NO	Yes
System Wide Data Dictionary	NO	Yes
Usable on System 37	NO	Yes
Upward/Downward compatible	NO	Yes

And This:

The RELATE™/3000 Relational Data Base Management System is a mature product with a fully developed, completely integrated set of fourth generation tool options. With five years of commercial acceptance and a proven record of customer satisfaction, the RELATE System provides you with the stability you need and the productivity you've only dreamed of.

Call Us Today for a Free Demonstration Or Product Information Packet

408/980-9898

CRI

CRI, INCORPORATED
Advanced Solutions for Management Problems
5333 Betsy Ross Drive, P.O. Box 58004
Santa Clara, CA 95052
(408) 980-9898

ENTER 245 ON READER CARD

ADVERTISERS INDEX

Reader Service Number		Page
100	Adager	26-27
191	Amos & Associates	96
196	Bering Industries	15
194	Bradford Business Systems	89
234	Bradmark Computer Systems	23
303	Cardamation Co. Inc.	97
216	Carolian Systems	91
103	Cognos Insert	
140	CompuChange	96
304	CompuChange	97
229	Computech Systems Corp.	87
136	Computer Media	96
230	Conam Corp.	83
301	Convertype	96
244	Copley Systems	73
186	Crisis Computer	96
211	CSU Industries	97
209	Data Systems For Industry	75
207	DBM Enterprises	97
176	Dynamic Information Systems Corp.	60-61
234	DTACK Systems	85
102	EMC Corporation	77
233	EMT Associates, Inc.	47
243	GFK	71
117	Graphicus	6
226	Hewlett-Packard (Networking)	38-39
200	Hewlett-Packard (Mass Storage)	8-9
242	Hewlett-Packard (Remarketing)	67
236	Holland House	40
162	Hypoint Technology	88
300	Hypoint Technology	96
305	Hypoint Technology	97
219	IMACS Systems	11
220	Indigo Software	31
235	Info Craft Inc.	85
198	Infocentre Ltd.	21
214	IOTECH	90
115	Martinsound	43
223	MicroTek Systems International	65
184	Norco	96
103	NSD, Inc.	3
231	Operations Control Systems	83
113	OPT	4
119	P.C.M.	63
232	Performance Software Group	7
116	Persoft, Inc.	I.B.Cover
240	Proactive Systems Ltd.	12-13
224	Robelle Consulting Ltd.	81
307	SEC International	97
134	Spectron Technology Inc.	68
306	System Works	97
185	Technical & Scientific Application	96
125	Telamon	35
	Tymlabs Corp. Supplement	
215	Walker Richer & Quinn	I.B.Cover
229	Walker Richer & Quinn	79
302	Wordware	96
111	Zentec	I.F.Cover

[CALENDAR]

[SEPTEMBER]

20-25: Interex North American Conference of Hewlett-Packard Business Computer Users, Las Vegas, NV. Contact Interex Conference Department, 680 Almanor Avenue, Sunnyvale, CA 94086; (408) 738-4848.

20-25: Management Information Systems For Strategic Advantage, Philadelphia, PA. Contact Registrar, Office of Executive Education, 200 Vance Hall, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104; (215) 898-4560.

[OCTOBER]

4-6: CAP International's 1987 Small Business DataTrack Service Conference, Grand Hyatt Hotel, New York, NY. Contact Jean O'Toole, CAP International, One Snow Road, Marshfield, MA 02050; (617) 837-1341.

18-22: Interex Technical Conference, San Jose, CA. Contact Interex; (408) 738-4848.

22: GHRUG (Greater Houston Regional Users Group), Hobby Hilton, Houston, TX. Contact Phil Curry; (713) 331-6111, ext. 255.

22-23: MARUG (North/South Carolina, Virginia). Fall Quarterly Meeting, Ocean Dunes Hotel, Myrtle Beach, NC. Conference theme centers on financial management systems. Contact Stephen Day; (804) 569-4857.

28-30: The World Conference On Electronic Printing & Publishing, Lisner Auditorium, George Washington University, 21st & H Streets NW, Washington, DC. Contact Henry B. Freeman; (703) 739-5510.

[NOVEMBER]

2-4: SuperGroup Users Conference-East, Washington Hilton, Washington, D.C. Contact The Producers, Riverwalk, 360 Merrimack St., Lawrence, MA 01843; (617) 683-5622.

9-12: NCGA's Mapping & Geographic Information Systems '87, The San Diego Princess Hotel, San Diego, CA. Event covers entire spectrum of the field from automated mapping/facilities management (AM/FM), to geographic information systems (GIS) to energy mapping. Contact Bob Cramblitt; (703) 698-9600.

HP Professional

Professional Press, Inc.
921 Bethlehem Pike
Spring House, PA 19477
(215) 542-7008

BUSINESS OFFICE

Associate Publisher/ Marketing & Sales Leslie Ringe

Sales Manager Jeffrey Berman

Marketing & Sales Coordinator Cynthia Leitzel

REGIONAL OFFICES

NEW ENGLAND/NEW YORK/ EASTERN CANADA
Cynthia Davis
238 Bedford Street, Suite 3
Lexington, MA 02173
(617) 861-1994

MID-ATLANTIC/ INTERNATIONAL
Leslie Ringe
238 Bedford Street, Suite 3
Lexington, MA 02173
(617) 861-1994

MIDWEST/SOUTHEAST/ SOUTHWEST
Peter Senft
921 Bethlehem Pike
Spring House, PA 19477
(215) 542-7008

SOUTHERN CALIFORNIA
David Beardslee
10635 Riverside Drive
Toluca Lake, CA 91602
(818) 985-7008

NORTHERN CALIFORNIA/ PACIFIC NORTHWEST
A. G. Germano
715 El Camino Real
Suite 206
San Bruno, CA 94066
(415) 873-3368

ONTARIO, CANADA
Helen Marbach
921 Bethlehem Pike
Spring House, PA 19477
(215) 542-7008

The SmarTerm® Theory of Software Evolution
What *Really* Happened to the Dinosaurs?

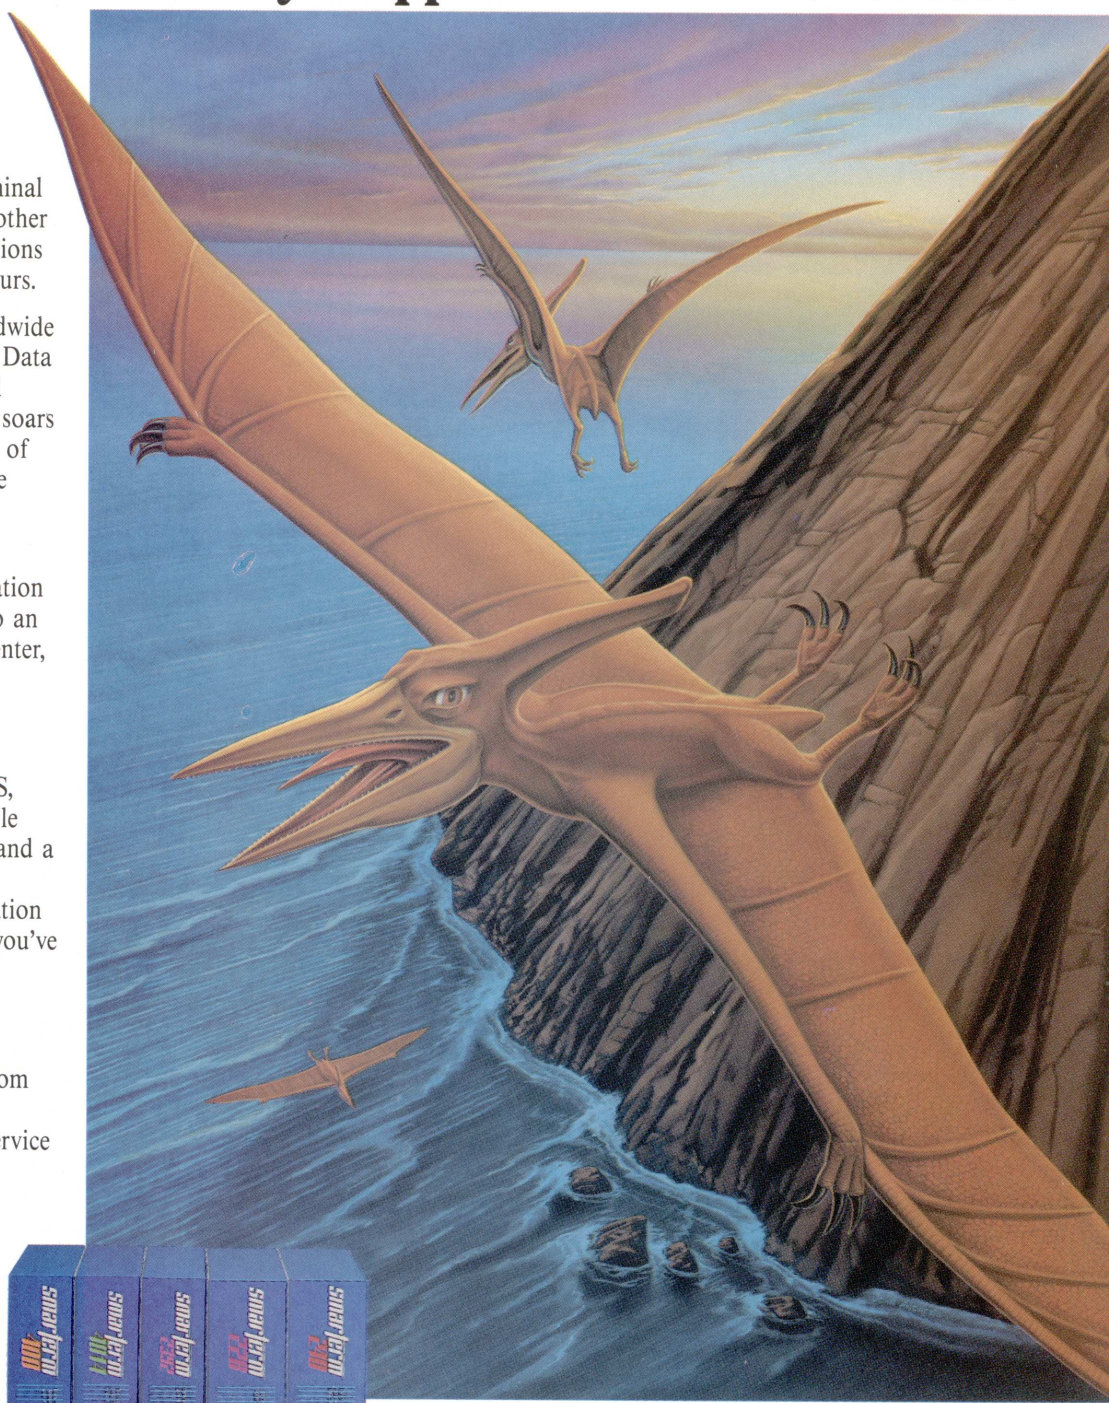

They were replaced with advanced new species.

After you've tried SmarTerm terminal emulation software, you'll think other PC-to-minicomputer communications links are as advanced as ...dinosaurs.

SmarTerm already is known worldwide as the leader in advanced DEC*, Data General* and Tektronix* terminal emulation software. Now Persoft soars to greater heights with emulation of another popular text terminal, the Hewlett-Packard* HP 2392.

SmarTerm 2392 offers you exact feature-for-feature terminal emulation and more. It connects your PC to an HP3000 Personal Productivity Center, another PC, and services like CompuServe* or Dow Jones News/Retrieval*.

Easy to use setup, Branch to DOS, Kermit and Xmodem error free file transfer, programmable softkeys, and a simple script language give you complete control. Add our dedication to quality technical support and you've got one "smart" communications software buy!

The SmarTerm family of terminal emulation products is available from your local dealer. Request more information through the reader service number below. Or contact:

Persoft, Inc.
465 Science Drive
Madison, WI 53711 USA
Phone (608) 273-6000
FAX (608) 273-8227
TELEX 759491

persoft®

ENTER 116 ON READER CARD

SmarTerm Terminal Emulation Software
. . . The Natural Selection

See us at Interex booth #307-406 and pick up your tickets for our hospitality party at Wet 'n Wild on September 24.

What keeps Reflection top of the line...

...is our bottom line.

SUPPORT

is as much a part of a terminal emulation package as the software itself.

That's why when you buy Reflection, you get **expert telephone support. Free of charge. Forever.**

There's extensive documentation, unlimited telephone support and an electronic bulletin board. Should a problem arise, a technical specialist will help you quickly trouble-shoot and resolve the matter.

The accuracy of Reflection's terminal emulation is legendary. Advanced features include command language, Local Area Network support, background processing, keyboard remapping, micro to host back-up, and several error-checking file transfer protocols. These features let your personal computer and host work as a team.

Reflection is available to connect a wide range of personal computers to Hewlett-Packard and Digital hosts. Call today for the connection that's right for you.

Your bottom line will reflect the results.

REFLECTION™
SERIES
SOFTWARE

See us at Interex Booths #307 - #40

ENTER 215 ON READER CARD

from Walker Richer & Quinn, Inc.
2825 Eastlake Avenue East, Seattle, WA 98102
1-800-872-2829